

## Release Note

# Rapier Layer 3 Gigabit Switch Software Release 2.1.0

Introduction .....	2
Rapier Switch Hardware Platform .....	3
Hardware Features .....	3
Uplink Modules .....	5
Network Service Modules (NSMs) .....	5
Port Interface Cards (PICs) .....	5
Power Supply .....	6
Configuring the Rapier Switch .....	6
Command Line Interface .....	6
Graphical User Interface .....	7
Logging In .....	8
Entering Commands .....	9
File Subsystem .....	9
Online Help .....	10
Configuration Scripts .....	11
Using the Built in Editor .....	12
Install Information .....	12
Downloading Releases and Patches into the Switch .....	14
Example: Install Software Upgrade for Rapier Switch .....	16
Interfaces .....	17
Layer 2 Switching .....	17
Switch Ports .....	18
Virtual LANs .....	25
The Switching Process .....	32
Quality of Service .....	36
Spanning Tree Protocol (STP) .....	38
IGMP Snooping .....	48
Triggers .....	50
Layer 3 Switching .....	51
Internet Protocol (IP) .....	51
Routing Information Protocol (RIP) .....	52
Novell IPX .....	52
AppleTalk .....	53
Resource Reservation Protocol (RSVP) .....	54
Layer 3 LAN/WAN Routing .....	55
SNMP and MIBs .....	56
Availability .....	58

# Introduction

---

Allied Telesyn announces the release of the Rapier 24, the first model in a new family of Layer 3 gigabit switches. This release note describes the new switch hardware platform, Layer 2 and Layer 3 switching features, and LAN/WAN multiprotocol routing features.

In addition to wire speed Layer 2 and Layer 3 IP switching, the Rapier family of switches implements the full AR router software suite from Allied Telesyn's AR series of routers., providing multiprotocol routing, IPsec, and the Nemesis stateful inspection firewall. For a complete description of the AR router software suite, see the AR Series Router Reference Manual (Document Number C613-03016-00 Rev B) at <http://www.alliedtelesyn.co.nz/support/rapier/>.

Before installing the Rapier 24, read the Rapier Switch Safety and Statutory Information (Document Number C613-02002-00 Rev A). Instructions for installing the Rapier switch are found in the Rapier Switch Quick Install Guide (Document Number C613-04017-00 Rev A). Both documents are supplied with the Rapier switch, or can be downloaded from <http://www.alliedtelesyn.co.nz/support/rapier/>.



**WARNING:** *Information in this release note is subject to change without notice and does not represent a commitment on the part of Allied Telesyn International. While every effort has been made to ensure that the information contained within the document and the features and changes described are accurate, Allied Telesyn International can not accept any type of liability for errors in, or omissions arising from the use of this information.*

---

Among the hardware features of the Rapier 24 are:

- Support for wire speed Layer 2 and Layer 3 IP switching.
- 24 autosensing 10/100 Ethernet ports.
- Two expansion bays for Gigabit Ethernet uplink modules.
- One Network Service Module bay for additional WAN interfaces.

The main software features of Software Release 2.1.0 are:

- Support for the Rapier 24 and its expansion options.
- Wire speed Layer 2 switching, including support for Virtual LANs.
- Wire speed Layer 3 IP switching.
- Layer 3 multiprotocol routing.

The documentation set for the Rapier switch includes:

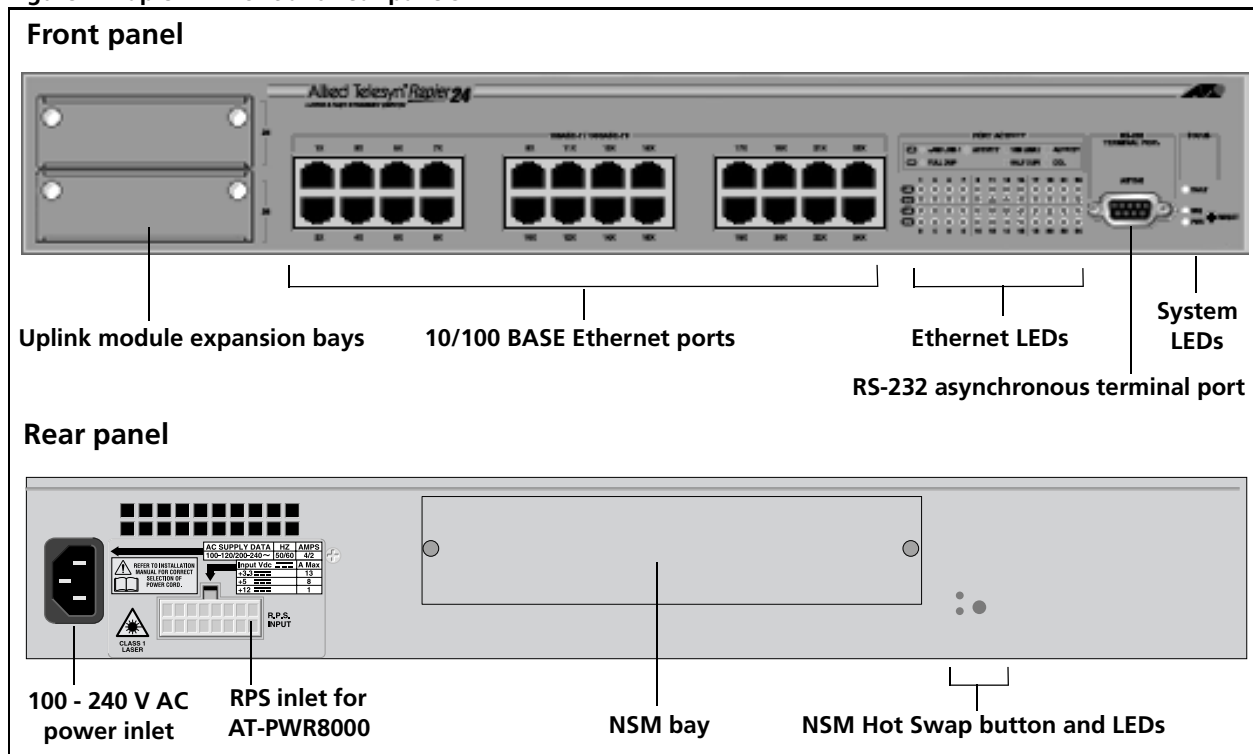
- Rapier Switch Safety and Statutory Information  
Document Number C613-02002-00 Rev A
- Rapier Switch Quick Install Guide  
Document Number C613-04017-00 Rev A
- Rapier Switch Documentation and Tools CD-ROM (in preparation), which includes the following PDF documents:
  - Rapier Switch Safety and Statutory Information
  - Rapier Switch Quick Install Guide
  - Rapier Switch Hardware Reference

- Rapier Switch Software Reference
- NSM Safety and Statutory Information
- NSM Quick Install Guide
- NSM Hardware Reference
- PIC Safety and Statutory Information
- PIC Quick Install Guide
- PIC Hardware Reference

## Rapier Switch Hardware Platform

The Rapier 24 combines high performance Layer 2 switching and Layer 3 IP switching with full multiprotocol routing capabilities in a single cost-effective package. Expansion options give upgradeable connectivity and support for multiple WAN services (Figure 1 on page 3).

Figure 1: Rapier 24 front and rear panels



## Hardware Features

The main features of the Rapier 24 switch are:

- High performance RISC-based architecture
- 1 MByte of EPROM
- 32 MBytes SDRAM
- 6 MBytes of FLASH memory
- 128 KBytes Non-Volatile Storage (battery backed SRAM)
- 24 x 10/100 Mbps autosensing Ethernet LAN ports

- 1 x RS-232 asynchronous serial port for switch management, with RS-232 DB9 cable for connection to terminal or PC.
- 2 x Uplink bays supporting Gigabit Ethernet uplink modules
- 1 x NSM (*Network Service Module*) bay
- A high performance 32-bit PAC slot for PCI accelerator card
- 110-240V AC power supply and optional redundant power supply (RPS)
- Support for the full AR switching and routing software suite
- 1.5U rack mounting
- LEDs indicating port activity and system status (Table 1).

**Table 1: Rapier 24 front panel LEDs**

LED	State	Function
<b>STATUS</b>		These LEDs indicate the state of the switch.
PWR	Green	The switch is receiving power and the voltage is in the acceptable range.
RPS	Green	A redundant power supply is connected to the switch and will provide power if the mains power fails or is disconnected.
Fault	Off	Normal operation.
	Flashing red	The switch is booting, running diagnostic tests, writing messages to FLASH memory, or transferring files using XMODEM.
	Red	The switch or management software is malfunctioning.
<b>PORT ACTIVITY</b>		These LEDs indicate the state of the switch ports.
L/A		Link/Activity
	Green	A 100 Mbps link is open.
	Flashing green	100 Mps activity is occurring.
	Amber	A 10 Mbps link is open.
	Flashing amber	10 Mps activity is occurring.
D/C		Duplex/collision
	Green	The port is operating at full-duplex.
	Amber	The port is operating at half-duplex.
	Flashing amber	Collisions are occurring on the line.

**Table 2: Rapier 24 rear panel LEDs**

LED	State	Function
<b>NSM</b>		These LEDs give indications about an NSM installed in the switch.
Swap	Green	Lit when the NSM is powered down and may be hot swapped. Only lit if the software also supports hot swapping.
In Use	Green	Lit when an NSM is installed. If the software supports hot swapping, it indicates that the NSM is powered up and may not be swapped.



---

*Note that this software release does not support hot swapping, so the switch must be powered down before an NSM is installed or removed. In future software releases, the Hot Swap button will be used to power down the NSM bay to allow hot swapping.*

---

## Uplink Modules

Each of the two uplink module expansion bays can support an optional Gigabit uplink module. The first uplink modules available are:

- AT-A35SX/SC, 1-port 1000BASE-SX (SC fibre connector)
- AT-A35LX/SC, 1-port 1000BASE-LX (SC fibre connector)

## Network Service Modules (NSMs)

The NSM bay accommodates *Network Service Modules* (NSMs) designed to support high speed LAN/WAN technologies. The NSM uses a 32Mhz 32-bit PCI style bus for high speed data applications. The first NSM to be released is the AT-AR040 4-PIC NSM, which has 4 PIC bays for installing Port Interface Cards (PICs). NSMs can be used interchangeably with the AT-AR740 router.

## Port Interface Cards (PICs)

The four PIC bays in Network Service Module AT-AR040 accommodate combinations of the following PIC slide-in interface cards:

- AR020 PRI E1/T1 PIC, 1 Primary Rate E1/T1 port
- AR021(S) BRI-S/T PIC, 1 Basic Rate ISDN S/T port
- AR021(U) BRI-U PIC, 1 Basic Rate ISDN U port
- AR023 SYN PIC, 1 Synchronous port with universal 50-way AMPLIMITE connector
- AR024 ASYN4 PIC, 4 Asynchronous ports with RJ45 connectors.

Most combinations of the PICs in “*Port Interface Cards (PICs)*” on page 5 can be installed in the PIC bays in the AT-AR040 NSM. Note the following limitations:

- Up to two PRI E1/T1 PIC cards (AT-AR020) can be installed in the NSM. If two PRI E1/T1 PICs are installed, the one must be in the lower row (bay 0 or 1), and the other must be in the upper row (bay 2 or 3).
- If a AT-AR020 PRI E1/T1 PIC is installed in one of the rows in the NSM and operating in E1 mode, then this row cannot also have an AT-AR021(S) BRI-S/T PIC or an AT-AR021(U) BRI-U PICU PIC installed.

The AT-AR040 NSM should be installed the switch, to give it mechanical and electrostatic protection, before installing PICs in the NSM.



---

*Install a PIC in the lowest numbered NSM PIC bay first, to avoid reallocating interface numbers when another PIC is installed.*

---

Changing the PICs in the lower numbered NSM bays when PICs of the same interface type are still installed in the higher numbered PIC bays will change the interface numbering, and therefore require changes to the software

configuration. PICs can be used interchangeably with Allied Telesyn's AR700 Series routers.

## Power Supply

The Rapier 24 has an AC power supply, which adapts to any power supply in the range of 110-240 VAC 50-60 Hertz input. It also has an inlet for connection to an optional redundant power supply (RPS) unit, AT-PWR8000. Future releases of the software will include power supply monitoring.

## Configuring the Rapier Switch

The Rapier is supplied with default settings which allow it to operate immediately as a switch, without any configuration. To take advantage of the full range of advanced Layer 2 switching features, the switch configuration must be changed. Layer 3 routing capabilities may also require detailed configuration. The switch has both a Command Line Interface (CLI) and a Graphical User Interface (GUI) for configuration and management.

## Command Line Interface

To use the command line interface (CLI) for configuring the switch, the first thing you need to do after physically installing the switch is to start a terminal session to access the switch (see Table 3 and the Rapier Switch Quick Install Guide). You will then be able to enter commands from this document and from the AR Series Router Reference Manual for Software Release 1.8 and 1.9.

To start a terminal session, do one of the following:

- Connect a VT100-compatible terminal to the RS-232 Terminal Port, set the communications parameters on the terminal (Table 3 on page 6), and press [Enter] a few times until the switch's login prompt appears; *or*
- Connect the COM port of a PC running terminal emulation software such as Windows Terminal or HyperTerminal to the RS-232 Terminal Port, set the communications parameters on the terminal emulation software (Table 3 on page 6), and press [Enter] a few times until the switch's login prompt appears; *or*
- Telnet to the switch from an IP host

**Table 3: Parameters for terminal communication**

Parameter	Value
Baud rate	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	Hardware

## Graphical User Interface

The switch includes a built-in web browser based GUI for configuring and monitoring Layer 2 switching features.

To enable the GUI, an IP address must be assigned to a switch interface. Connect a terminal or a PC running terminal emulation software (for instance Windows Terminal or HyperTerminal) to the RS-232 Terminal Port, and log in to the manager account.

Add an IP interface over the default VLAN (vlan1) and assign it an IP address (e.g. 192.168.1.1), using the command:

```
ADD IP INTERFACE=vlan1 IPADDRESS=192.168.1.1
```

Enable the GUI with the command:

```
ENABLE GUI
```

Point your web browser at the IP address you have assigned to the switch, for example, <http://192.168.1.1>. The authentication window appears (Figure 2). The “Enter Network Password dialog” allows only authorised users with the correct password to access the router. (The appearance of this dialog may differ depending on the browser you use.)

**Figure 2: Enter Network Password dialog**



- Enter the username *manager* and the password *friend*. The GUI Main Screen opens (Figure 3).

Figure 3: Rapier 24 Welcome page



While using the Rapier GUI, use the buttons on the pages to navigate, not the browser's Back and Forward buttons, to ensure that configuration information is stored correctly.

If you have not yet changed the initial manager password, we recommend that you do this now. Make sure you remember the new password, as there is no way to retrieve it if it is lost.

To leave the Rapier GUI, click the Exit button on the Welcome page.

## Logging In

A user accessing the switch from a terminal or PC connected to the front panel RS-232 terminal port (asyn0), or via a Telnet connection, must enter a login name and password to gain access to the command prompt. When the switch is supplied, it has a *manager* account with an initial password *friend*. Enter your login name at the login prompt:

```
login: manager
```

Enter the password at the password prompt:

```
password: friend
```

This password should be changed to prevent unauthorised access to the switch, using the command:

```
SET PASSWORD
```

Make sure you remember the new password you create, as a lost password cannot be retrieved, and would mean losing access for configuring and monitoring the switch.

The command processor supports three levels of privilege, USER, MANAGER, and SECURITY OFFICER, distinguished by the prompt displayed by the

command processor when it is ready to receive commands. A USER level prompt looks like:

>

while a MANAGER prompt looks like:

Manager >

and a SECURITY OFFICER prompt looks like:

SecOff >

See the *Operations Chapter* of the *AR Series Router Reference Manual* at <http://www.alliedtelesyn.co.nz/support/rapier/> for more information about creating new accounts with user, manager and security officer privileges.

## Entering Commands

The switch is controlled with the commands in this document and in the AR Series Router Reference Manual for Software Releases 1.8 and 1.9. While the keywords in commands are not case sensitive, the values entered for some parameters are. The switch supports command line editing and recall (Table 4 on page 9).

**Table 4: Command line editing functions and keystrokes**

Function	VT100-compatible Keystroke
Move cursor within command line	←, →
Delete character to left of cursor	[Delete] or [Backspace]
Toggle between insert/overstrike	[Ctrl/O]
Clear command line	[Ctrl/U]
Recall previous command	↑ or [Ctrl/B]
Recall next command	↓ or [Ctrl/F]
Display command history	[Ctrl/C] or SHOW ASYN HISTORY
Clear command history	RESET ASYN HISTORY
Recall matching command	[Tab] or [Ctrl/I]

## File Subsystem

FLASH memory is structured like a file subsystem. Files can be saved, renamed, listed and deleted. Release files, online help files, configuration scripts and other scripts are all stored as files in FLASH memory. Names must have DOS format, with a filename of up to eight characters and an extension of three characters.

To display the files in FLASH, use the command:

SHOW FILE

Figure 4: Example output from the SHOW FILE command.

Filename	Device	Size	Created	Locks
-----	-----	-----	-----	-----
1mac.scp	flash	527	08-Nov-2000 12:46:00	0
86s-210.rez	flash	1690736	14-Sep-2000 14:11:56	0
config.scp	flash	64	10-Nov-2000 23:26:31	0
hdroute.scp	flash	374	08-Nov-2000 12:46:00	0
loadup.scp	flash	173	20-Nov-2000 07:03:30	0
loadup1.scp	flash	224	14-Nov-2000 14:11:56	0
quick.scp	flash	2036	08-Nov-2000 12:46:00	0
release.lic	flash	32	08-Nov-2000 12:46:00	0
sleep.scp	flash	189	08-Nov-2000 12:46:00	0
test.cfg	flash	1698	09-Nov-2000 10:39:42	0
-----	-----	-----	-----	-----

The switch automatically compacts FLASH memory when a maximum threshold of deleted files is reached. Compaction frees space for new files by discarding garbage. A message will appear when FLASH compaction has been activated. Another message appears when FLASH compaction is complete.



*While FLASH is compacting, do not restart the switch or use any commands that affect the FLASH file subsystem. Do not restart the switch, or create, edit, load, rename or delete any files until a message confirms that FLASH file compaction is completed. Interrupting flash compaction may result in damage to files.*

## Online Help

Online help is available for all switch commands. Typing a question mark “?” at the end of a partially completed command displays a list of the parameters that may follow the current command line, with the minimum abbreviations in uppercase letters. The current command line is then re-displayed, ready for further input.

An online help facility provides more detailed help information via the command:

```
HELP [topic]
```

If a topic is not specified, a list of available topics is displayed. The HELP command displays information from the system help file stored in FLASH memory. The help file used by the HELP command must be defined using the command:

```
SET HELP=helpfile
```

The current help file and other system information can be displayed with the command:

```
SHOW SYSTEM
```

Figure 5: Example of output from the SHOW SYSTEM command

```

Switch System Status                               Time 14:29:17 Date 12-Sep-2000.
Board      ID   Bay Board Name                      Rev   Serial number
-----
Base       86   AT-RP24 Rapier 24                          P2-1  49867449
-----
Memory -   DRAM : 32768 kB   FLASH : 6144 kB
-----
SysDescription
CentreCOM AT-RP24 Rapier 24 version 2.1.0-00 04-Sep-2000
SysContact

SysLocation

SysName

SysUpTime
30262 ( 00:05:02 )
Software Version: 2.1.0-00 04-Sep-2000
Release Version : 2.1.0-00 04-Sep-2000
Release built   : Sep 12 2000 at 14:28:59
Patch Installed : NONE
Territory       : usa
Help File       : help.hlp

Main PSU        : On           Main Fan         : On
RPS Monitor     : On           RPS Connected : Yes
RPS PSU         : On           RPS Fan       : On

Boot configuration file: vts.cfg (exists)
Current configuration: vts.cfg
Security Mode    : Disabled

Warning (248283): No patches found.

```

## Configuration Scripts

At boot the switch executes the commands in the boot script to configure the switch. A boot script is a sequence of standard commands that the switch executes at start-up. The default boot script is called `boot.cfg`, but an alternative script file can be defined as the boot script using the command:

```
SET CONFIG=filename
```

Subsequent commands entered from the command line or executed from a script affect only the dynamic configuration in memory, which is not retained over a power cycle. Changes are not automatically stored in nonvolatile memory. When the switch is restarted the configuration will be restored to that defined by the boot script, or if the switch was restarted using the RESTART command, any script specified in the RESTART command.

To ensure that any configuration changes made after boot are retained across a restart or power cycle, the modified configuration must be saved as a script file, using the command:

```
CREATE CONFIG=filename
```



*The CREATE CONFIG command writes the MD5 digest, not the clear text, of passwords in commands to the configuration file. When a configuration script is*

*executed the command processor can determine whether the password value is clear text or an MD5 digest.*

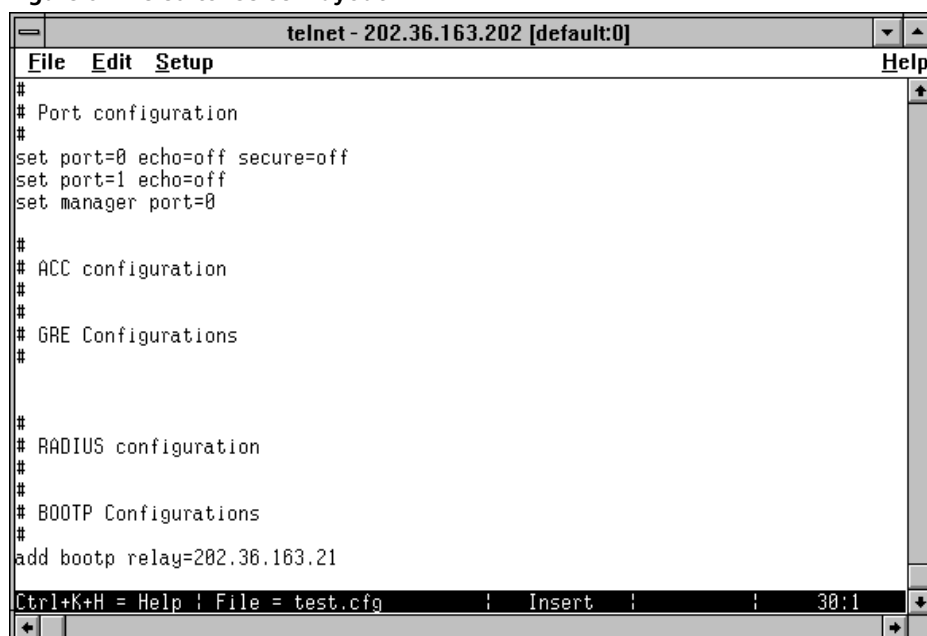
## Using the Built in Editor

The switch has a built-in full-screen text editor for editing script files stored on the switch file subsystem. Scripts can be run manually, or run when a trigger automatically activates on some specified events in the switch. See “Triggers” on page 50, and the *Trigger Facility* chapter in the *AR Series Router Reference Manual* at <http://www.alliedtelesyn.co.nz/support/rapier/>. The editor is invoked with the command:

```
EDIT [filename]
```

The file name is optional as a file can be loaded, or a new file can be created from within the editor itself (Figure 6 on page 12).

**Figure 6: The editor screen layout.**



The editor uses VT100 command sequences and should only be used with a VT100-compatible terminal, terminal emulation program or Telnet client.

To display Help at any time while in the editor press [Ctrl/K,H]; that is, hold down the Ctrl key and press in turn the K key then the H key.

## Install Information

The INSTALL module is responsible for maintaining install information and loading the correct install at boot. A release is a binary file containing the code executed by the switches CPU. There may also be a patch file, and additional binary file that modifies the original release file. An *install* is a record identifying a release and an optional patch. Three installs are maintained by the INSTALL module, *temporary*, *preferred* and *default*.

The default install is the install of last resort. The release for the default install can not be changed by the manager and is always the EPROM release. The patch for the default install may be set by the manager.

The temporary and preferred installs are completely configurable. Both the release and an associated patch may be set. The release may be EPROM or a release stored in FLASH.

The three different installs are required to handle the following situations:

- A default install is required to handle the case when only the EPROM release is present.
- A temporary install is required to allow a release and/or patch to be loaded once only, in case it causes a switch crash.
- A preferred install is required because the default install can not be anything other than the EPROM.

The install information is inspected in a strict order. The temporary install is inspected first. If this install information is present, the temporary install is loaded. At the same time, the temporary install information is deleted. This ensures that if the switch reboots immediately as the result of a fatal condition caused by the temporary install, the temporary install will not be loaded a second time.

If there is no temporary install defined, or the install information is invalid, the preferred install is inspected. If present, this install is loaded. The preferred install information is never deleted.

If neither temporary nor preferred installs are present, the default install is used. The default install will always be present in the switch, because if, for some reason, it is not, the INSTALL module will restore it.




---

*The preferred install should not be set up with an untested release or patch. It is advisable to install new releases or patches as the temporary install, and when the switch boots correctly, to then set up the preferred install with the new release or patch.*

---

To change the install information in the switch, use the command:

```
SET INSTALL={TEMPORARY|PREFERRED|DEFAULT}
    [RELEASE={release-name|EPROM}] [PATCH={patch-name}]
```

The INSTALL parameter specifies which install is to be set. The INSTALL module is responsible for maintaining install information and loading the correct install at boot. An *install* is a record identifying a release and an optional patch. Three installs are maintained by the INSTALL module, *temporary*, *preferred* and *default*.

The default install is the install of last resort. The release for the default install can not be changed by the manager and is always the EPROM release. The patch for the default install may be set by the manager.

The temporary and preferred installs are completely configurable. Both the release and an associated patch may be set. The release may be EPROM or a release stored in FFS.

The RELEASE parameter specifies the release file for this install. The release file is either a file name of the form `device:filename.ext` for files in the file subsystem, or EPROM, to indicate the EPROM release. The default value for the device field is FLASH.

The PATCH parameter specifies the patch file for this install, and is a file name of the form `device:filename.ext`. The patch file is stored in FLASH. The default value for the device field is FLASH. If the patch name is not given, the patch file information for a given install is removed and only the release file will be loaded for the install.

A patch file can not be set up for an install unless a release file is already set up, or a release file is specified in the same command. This stops the inadvertent setting of an install to be just a patch file. When the switch reboots in such a case the particular install is ignored, which may have undesirable effects on the switch operation.



*For security reasons this command will only be accepted if the user has SECURITY OFFICER privilege.*

To delete a particular install (except the default install) use the command:

```
DELETE INSTALL
```

To display the current install information, including which install is currently running in the switch, and how the install information was checked at the last reboot, use the command:

```
SHOW INSTALL
```

**Figure 7: Example output from the SHOW INSTALL command.**

Install	Release	Patch	Dmp
Temporary	-	-	-
Preferred	flash:86s-210.rez	-	-
Default	EPROM (8-1.6.0)	-	-
Current install			
Preferred	flash:8d-181.rez	-	-
Install history			
No Temporary install selected			
Preferred install selected			
Preferred release successfully installed			
Preferred patch successfully installed			

## Downloading Releases and Patches into the Switch

The LOADER module is responsible for loading and storing releases, patches and other files into FLASH. The LOADER module uses the *Trivial File Transfer Protocol* (TFTP), *Hypertext Transfer Protocol* (HTTP) or ZMODEM over an asynchronous port, to retrieve files from a network host. The FFS module is used to create, write and destroy release and patch files.

The loader can be configured with the command:

```
SET LOADER [DELAY=delay|DEFAULT]
           [DESTINATION={FLASH|DEFAULT}] [FILE=filename]
           [HTTPPROXY={hostname|ipadd|DEFAULT}]
           [METHOD={HTTP|TFTP|WEB|WWW|ZMODEM|NONE|DEFAULT}]
           [ASYN=port|DEFAULT] [PROXYPORT=1..65535|DEFAULT]
           [SERVER={hostname|ipadd|DEFAULT}]
```

This command sets default values for the name of the file to load, the network host to load it from, and the memory location in which to store the file. These default values can be overridden when the load actually takes place. A time delay between initiating a load and the start of the load can also be configured.

The DELAY parameter specifies the delay, in seconds, between initiating the file download and the download actually starting. This feature is provided to allow reconfiguration of ports and devices after initiating the download. For example, a manager may be at a remote site with a single PC which is to act as both the access device to the switch and the TFTP server. By specifying a delay, the manager has time to reconfigure the PC from terminal emulation mode to TFTP server mode before the download starts. The DELAY parameter is optional. If DEFAULT is specified, this parameter is set to the factory default, which is no delay.

The DESTINATION parameter specifies where the file will be stored. If FLASH is specified, the file is stored in the FLASH File System (FFS) on the switch. If DEFAULT is specified, this parameter is set to the factory default, FLASH.

The FILE parameter specifies the name of the file, in the syntax of the server from which the file will be downloaded. The FILE parameter is a full path name rather than just a file name. The only restriction is that the last part of the parameter must be a valid file name for the LOADER module. When METHOD is set to TFTP, HTTP, ZMODEM or NONE, valid file names are of the form *filename.ext* where *filename* is one to eight characters in length and *ext* is three characters in length. The following are examples of valid file names for methods TFTP, ZMODEM or NONE:

```
\user\public\filename.ext ; UNIX or DOS server
[network.cfg]filename.ext ; DEC VAX server
```

Note that, starting at the end of the file name and working backwards, the first character not valid in file names delimits a valid file name for the switch. If the slash at the beginning of the path is omitted in this command, the LOAD command adds it. The following are examples of valid file names for method HTTP:

```
/path/filename.ext
path/filename.ext
```

The HTTPPROXY parameter specifies the proxy server used to handle HTTP requests. Either the IP address or the fully qualified domain name of the proxy server may be specified. If a domain name is specified, the switch will perform a DNS lookup to resolve the name. If DEFAULT is specified, this parameter is set to the factory default, which has no value set for HTTPPROXY, clearing any value previously set as default.

The METHOD parameter specifies the method to use when downloading the file. If HTTP is specified, HTTP is used to download the file. The options WEB and WWW are synonyms for HTTP. If TFTP is specified, TFTP is used to download the file. If ZMODEM is specified, the ZMODEM protocol is used to download the file. If ZMODEM is specified, the PORT parameter must be specified, unless it has been set with the SET LOADER command. If NONE is specified, only text files can be downloaded and all input received via the port will be directed to the specified file on the switch's file subsystem. The file

transfer is terminated by the first control character received that is not a CR or LF character. The FILE parameter is not valid when METHOD is set to ZMODEM. The PORT parameter is not valid when METHOD is set to HTTP, WEB, WWW, TFTP or NONE. If DEFAULT is specified, this parameter is set to the factory default, which is TFTP.

The ASYN parameter specifies the asynchronous port via which the file will be downloaded, when the METHOD parameter is set to ZMODEM or NONE. If METHOD is set to ZMODEM or NONE, the PORT parameter is required unless it has been set with the SET LOADER command. If DEFAULT is specified, this parameter is set to the factory default, which is no PORT set, clearing any value previously set as default.

The PROXYPORT parameter specifies the port on a proxy server. The PROXYPORT parameter is only valid if METHOD is HTTP and HTTPPROXY is specified. If DEFAULT is specified, this parameter is set to the factory default, which is 80.

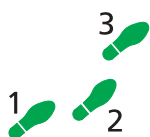
The SERVER parameter specifies the IP address or the host name (a fully qualified domain name) of the TFTP server or HTTP server from which the file is loaded. If a host name is specified, a DNS lookup is used to translate this to an IP address. The SET IP NAMESERVER command can be used to define name servers. The PING command can be used to verify that the switch can communicate with the server via IP. The SERVER parameter is not used when METHOD is set to ZMODEM or NONE. The following are examples of valid server names when METHOD is set to HTTP:

```
host.company.com
192.168.3.4
```

If DEFAULT is specified, this parameter is set to the factory default, which has no value set for SERVER, clearing any value previously set as default.

## Example: Install Software Upgrade for Rapier Switch

This example downloads a compressed release from the Rapier Support site to the switch's FLASH memory using HTTP.



### To install a compressed release:

#### 1. Download the release files to the router.

The release file is downloaded to the switch with the command:

```
LOAD METHOD=HTTP DESTINATION=FLASH
FILE=/support/rapier/downloads/86s-210.rez
SERVER=www.alliedtelesyn.co.nz HTTPPROXY=proxy-address
PROXYPORT=proxy-port
```

where *proxy-address* is the fully qualified domain name (e.g. proxy.mycompany.com) or IP address (e.g. 192.168.1.1) of the proxy server, and *proxy-port* is the port number of the proxy port on the proxy server. If access from the switch to the world wide web is not via a proxy server, the HTTPPROXY and PROXYPORT parameters should be omitted.

The process of downloading a release file can take some time, even if the router and the HTTP server are connected by high speed links. An indicative time for downloading a release over Ethernet is 5 to 10 minutes. The progress of the download can be monitored with the command:

```
SHOW LOAD
```

When the download has completed, the presence of the files in FLASH can be displayed with the command:

```
SHOW FILE
```

This shows the file 86s-210.rez is present.

## 2. Test the release.

The release can now be tested, using the command:

```
SET INSTALL=TEMPORARY RELEASE=86s-210.REZ
```

The install information can be checked with the command:

```
SHOW INSTALL
```

The switch is then rebooted, and the install is checked again. This display should indicate, in the install history, that the temporary install was loaded.

## 3. Make the release the default (permanent) release.

If the switch operates correctly with the new release, the release may be made permanent with the command:

```
SET INSTALL=PREFERRED RELEASE=86s-210.REZ
```

Every time the router reboots from now on, the new release will be loaded from FLASH.

Other load methods are described in the *Operations* chapter in the *AR Series Router Reference Manual* at <http://www.alliedtelesyn.co.nz/support/rapier/>.

# Interfaces

Interface naming for asynchronous interfaces and Ethernet interfaces on the switch differs from that described in the AR Series Router Reference Manuals for Software Releases 1.8 and 1.9. The asynchronous interfaces called 'portn' in the AR router manuals (for instance 'INTERFACE=port0') are referred to as 'asynn' on the Rapier switch (for instance 'asyn0'). Ethernet ports, referred to as 'ethn' in the AR router manuals (for instance 'INTERFACE=eth0'), are numbered from one, and simply referred to in most switch commands by their number (for instance 'PORT=1'). The Testing Facility refers to these interfaces as 'portn' (for instance 'port1').

Interfaces can be configured to VLANs for IP, IPX and Appletalk routing protocols in the same way that other interfaces are created for other interface types. Concatenate VLAN with the VID of the VLAN giving VLANn, for instance:

```
INTERFACE=VLAN3
```

# Layer 2 Switching

This section describes the Layer 2 switching features on the Rapier switch, and how to configure them.

The main Layer 2 features of Software Release 2.1.0 are:

- High performance, non-blocking, wire-speed Layer 2 switching.

- Packet Forwarding at wire speed.
- Store and Forward switching mode.
- Autonegotiation of link speed and duplex mode for 10/100 Mbps speed on all 100BASE TX ports.
- Autonegotiation of duplex mode for 10/100 and gigabit Ethernet ports.
- Automatic, configurable MAC address learning and ageing, supporting up to 8191 MAC addresses per switch.
- Switch Filtering.
- Flow Control.
- Broadcast Storm Protection.
- Spanning Tree Protocol.
- Up to 62 Virtual LANs defined by port membership ("*Virtual LANs*" on page 25).
- Priority tagging to support four QOS egress queues.
- Port trunking to spread traffic over several links.
- Port mirroring.
- IGMP (Internet Group Management Protocol) snooping.

## Switch Ports

Each port is uniquely identified by a port number. The switch supports a number of features at the physical level that allow it to be connected in a variety of physical networks. This physical layer (layer 1) versatility includes:

- Enabling and disabling of Ethernet ports.
- Auto negotiation of port speed and duplex mode for all 10/100 Ethernet ports.
- Manual setting of port speed and duplex mode for all 10/100 Ethernet ports.
- Setting flow control parameters for all ports.
- Link up and link down triggers.
- Port trunking.
- Packet storm protection.
- Port mirroring.
- Support for SNMP management

## Enabling and disabling Ethernet ports

An Ethernet port that is enabled is available for packet reception and transmission. Its administrative status in the Interfaces MIB is UP. Conversely, an Ethernet port that is disabled is not available for packet reception and transmission. It will not send or receive frames, or participate in spanning tree negotiation. Its administrative status in the Interfaces MIB is DOWN. Every Ethernet port on the switch is enabled by default, and can be enabled and disabled using the commands:

```
ENABLE SWITCH PORT=port-list
```

```
DISABLE SWITCH PORT=port-list
```

where:

- *port-list* is a port number, a range of port numbers (specified as n-m), or a comma separated list of port numbers and/or ranges. Port numbers start at 1 and end at m, where m is the highest numbered switch Ethernet port. (including uplink ports).

To display information about the settings of the switch port, use the command:

```
SHOW SWITCH PORT=port-list
```

**Figure 8: Example output from the SHOW SWITCH PORT command.**

```
Switch port Information
-----
Port ..... 1
Link state ..... Up
UpTime ..... 11210 s
Port Media Type ..... ISO8802-3 CSMACD
Configured speed/duplex ..... Autonegotiate
Actual speed/duplex ..... 100 Mbps, full duplex
Acceptable Frames Type ..... Admit All Frames
Broadcast rate limit ..... 1000/s
Multicast rate limit ..... -
DLF rate limit ..... -
Learn limit ..... 20
Lock action ..... Discard
Current learned, lock state ... 15, not locked
Mirroring ..... None
Is this port mirror port ..... No
Port VLAN Identifier ..... 42
Port-based VLAN ..... accounting (42)
Tagged VLAN ..... marketing (87)
                        sales (321)
Ingress Filtering ..... OFF
STP ..... company

Port ..... 2
Link state ..... Up
UpTime ..... 1545 s
Port Media Type ..... ISO8802-3 CSMACD
Configured speed/duplex ..... 10 Mbps, half duplex
Actual speed/duplex ..... 10 Mbps, half duplex
Acceptable Frames Type ..... Admit All Frames
Broadcast rate limit ..... 1000/s
Multicast rate limit ..... -
DLF rate limit ..... -
Learn limit ..... None
Lock action ..... None
Current learned, lock state ... 15, not locked
Mirroring ..... Tx, to port 22
Is this port mirror port ..... No
Port VLAN Identifier ..... 1
Port-based VLAN ..... default (1)
Tagged VLAN(s) ..... -
Ingress Filtering ..... Off
STP ..... company
-----
```

**Table 5: Parameters in the output of the SHOW SWITCH PORT command**

Parameter	Meaning
Port	The number of the switch port.
Link state	The link state of the port, one of "Up" or "Down".
Uptime	The count in seconds of the elapsed time since the port was last reset or initialised.
Port Media Type	The MAC entity type as defined in the MIB object ifType.
Configured speed/duplex	The port speed and duplex mode configured for this port. One of "Autonegotiate" or a combination of a speed (one of "10 Mbps", "100 Mbps" or "1000 Mbps") and a duplex mode (one of "half duplex" or "full duplex").
Actual speed/duplex	The port speed and duplex mode that this port is actually running at. A combination of a speed (one of "10 Mbps", "100 Mbps" or "1000 Mbps") and a duplex mode (one of "half duplex" or "full duplex").
Acceptable Frames Type	The value of the Acceptable Frames Type parameter, one of: "Admit All Frames" or "Admit Only VLAN-tagged Frames".
Broadcast rate limit	The limit of the rate of reception of broadcast frames for this port, in frames per second.
Multicast cast rate limit	The limit of the rate of reception of multicast frames for this port, in frames per second.
DLF rate limit	The limit of the rate of reception of DLF (destination lookup failure) frames for this port, in frames per second.
Learn limit	The number of MAC addresses that may be learned for this port. Once the limit is reached, the port is locked against any new MAC addresses. One of "None" or a number from 1 to 256.
Lock action	The action taken on this port when a frame is received from an unknown MAC address when the port is locked. One of "None", "Discard", "Trap" or "Disable".
Current learned, lock state	The number of MAC addresses currently learned on this port and the state of locking for this port. The lock state is one of "not locked", "locked by limit" or "locked by command".
Mirroring	The traffic mirroring for traffic in and out of this port. One of "None", "Rx" (for traffic received by this port), "Tx" (for traffic sent on this port) or "Both". The port to which mirrored frames are being sent is also displayed.
Is this port mirror port	Whether or not this port is a mirror port. One of "No" or "Yes".
Port VLAN Identifier	The VLAN Identifier (VID) that may be associated with untagged or priority-tagged frames.
Port-based VLAN	The name and VLAN Identifier (VID) of the port-based VLAN to which the port belongs.
Tagged VLAN	The name and VLAN Identifier (VID) of the tagged VLAN(s), if any, to which the port belongs.
Ingress Filtering	The state of Ingress Filtering: one of "On" or "Off".
STP	The name of the STP to which the port belongs.

Resetting Ethernet ports at the hardware level discards all frames queued for reception or transmission on the port, and restarts autonegotiation of port speed and duplex mode. This clears any packets stuck in a queue, for instance after a broadcast storm, and may sometimes make a non-operational port operational again. Ports are reset using the command:

```
RESET SWITCH PORT=port-list
```

## Auto negotiation of port speed and duplex mode

Each of the 10/100 Ethernet ports on the switch can operate at either 10Mbps or 100 Mb per second, in either full duplex or half duplex mode. In full duplex mode a port can transmit and receive data simultaneously, while in half duplex mode the port can either transmit or receive, but not at the same time. This versatility makes it possible to connect devices with different speeds and duplex modes to different ports on the switch. Such versatility also requires that each port on the switch know which speed and mode to use.

Autonegotiation allows the ports to adjust their speed and duplex mode to accommodate the devices connected to them. Each 10/100 Ethernet port can be either configured with a fixed speed and duplex mode, or configured to autonegotiate speed and duplex mode with a device connected to it to determine a speed and mode that will allow successful transmission. If another autonegotiating device is connected to the switch, they will negotiate the highest possible common speed and duplex mode (Table 6). Setting the port to a fixed speed and duplex mode allows it to support equipment that cannot autonegotiate. 10/100 Ethernet ports will autonegotiate by default when they are connected to a new device. To change this setting, use the command:

```
SET SWITCH PORT=port-list
SPEED={ AUTONEGOTIATE | 10MHALF | 10MFULL | 100MHALF | 100MFULL | 1000MHALF | 1000MFULL }
```

Autonegotiation can also be activated at any time after this, on any port that is set to autonegotiate by using the command:

```
ACTIVATE SWITCH PORT=port-list AUTONEGOTIATE
```

The AUTONEGOTIATE parameter specifies that the port is to activate the autonegotiation process. The port will begin to autonegotiate link speed and duplex mode.

On the first Rapier 24 switch to be released the Gigabit uplink ports always use 1000 Mbps full duplex mode, but these can also autonegotiate with peers in order to successfully pass the negotiation phase to get to successful operation. (This is a limitation on the B1 revision switch silicon, and may be lifted in later revisions.)

**Table 6: Autonegotiation preferences for Ethernet ports**

Preference	10/100 Ethernet ports	Gigabit uplink ports
Highest	100 Mbps Full duplex	1000Mbps Full duplex
	100 Mbps Half duplex	
	10Mbps Full duplex	
Lowest	10Mbps Half duplex	

## Flow control

Ethernet flow control is designed to cope with the situation where packets received on an Ethernet port are queued due to a limitation in available bandwidth on the port or ports out which those packets are switched. The

senders of traffic on ports which have too many packets in the input queue can be informed of the situation and made to restrict the flow of packets.

Flow control for Ethernet ports consists of two mechanisms. The mechanism for a given port is dependent on the duplex mode of the port. For full duplex ports, flow control is achieved by sending a special PAUSE MAC frame out the port, which tells the sending device not to send any more frames for a random period of time. This mechanism is defined in IEEE Standard 802.3. For half duplex ports, flow control is achieved by asserting the jamming signal on the port, a mechanism known as back pressure. The jamming signal causes collisions on the link, so that the sending device waits for a random period of time before sending any more frames.

The maximum size of the ingress queue depends on the model. For any given port on the switch, there will be a command to control whether either of the mechanisms for flow control will be employed. The following table illustrates this. This shows the action taken on a port of given duplex mode with the various combinations of flow control enabled, when the flow control state is entered.

Both kinds of flow control are enabled by default on every Ethernet port, and they can be enabled and disabled using the command:

```
ENABLE SWITCH PORT=port-list
FLOW={JAMMING | PAUSE} [, {PAUSE | JAMMING}]

DISABLE SWITCH PORT=port-list
FLOW={JAMMING | PAUSE} [, {PAUSE | JAMMING}]
```

The FLOW parameter specifies the types of flow control to be enabled for the port. One or both types may be enabled with this command. If JAMMING is specified, flow control for half duplex ports by asserting the jamming signal will be enabled. If PAUSE is specified, flow control for full duplex ports by sending PAUSE frames will be enabled. Both these forms of flow control are enabled by default.

Table 7 shows the effect of flow control parameter settings on the behaviour of ports when maximum ingress queues are exceeded.

**Table 7: Effect of flow control parameters on ports**

Flow control parameter settings	Half duplex	Full duplex
PAUSE disabled and jamming disabled	packets discarded	packets discarded
PAUSE disabled and jamming enabled	assert jamming	packets discarded
PAUSE enabled and jamming disabled	packets discarded	send PAUSE
PAUSE enabled and jamming enabled	assert jamming	send PAUSE

For the Rapier switch, jamming on all half duplex ports is set by a single parameter. The Rapier switch will not reach the jamming threshold, but will discard packets before reaching this limit.

## Port trunking

Port trunking, also known as port bundling or link aggregation, allows a number of ports to be configured to join together to make a single logical connection of higher bandwidth. This can be used where a higher performance link is required, and makes links even more reliable.

The Rapiert switch supports up to 6 trunk groups, of up to 8 10/100 Ethernet ports each. The two 2 gigabit Ethernet ports can also be grouped together to form a trunk group. Ports do not have to be contiguous. Port trunking can be used between any two Rapiert switches.

Port trunk groups are created and destroyed on the switch using the commands:

```
CREATE SWITCH TRUNK=trunk [SPEED={10M|100M|1000M}]
[SELECT={MACSRC|MACDEST|MACBOTH|IPSRC|IPDEST|IPBOTH}]
DESTROY SWITCH TRUNK=trunk
```

The members of a trunk group can be specified when it is created, and ports can be added to or removed from a trunk group using the commands:

```
ADD SWITCH TRUNK=trunk PORT=port-list
DELETE SWITCH TRUNK=trunk PORT=port-list
```

On the Rapiert switch, ports which are members of a trunk group must be configured for full duplex mode. When a port is added to a trunk group, the speed setting for the group overrides the speed setting previously configured for the port. The speed of the trunk group can either be specified when it is created, or set using the command:

```
SET SWITCH TRUNK=trunk SPEED={10M|100M|1000M}
[SELECT={MACSRC|MACDEST|MACBOTH|IPSRC|IPDEST|IPBOTH}]
```

The SELECT parameter optionally specifies the port selection criterion for the trunk group. Each packet to be sent on the trunk group is checked, using the selection criterion, and a port in the trunk group chosen down which to send that packet. If MACSRC is specified, the source MAC address is used. If MACDEST is specified, the destination MAC address is used. If MACBOTH is specified, both source and destination MAC addresses are used. If IPSRC is specified, the source IP address is used. If IPDEST is specified, the destination IP address is used. If IPBOTH is specified, both the source and destination IP addresses are used. The user of the switch should choose this parameter to try to spread out the load as evenly as possible on the trunk group. The default for this parameter is MACDEST.

To display information about trunks on the switch, use the command:

```
SHOW SWITCH TRUNK [=trunk]
```

**Figure 9: Example output from the SHOW SWITCH TRUNK command**

Switch trunk groups	
-----	
Trunk group name .....	Uplink
Speed .....	1000Mbps
Selection criterion .....	Destination MAC address
Ports .....	25,26
-----	

## Packet storm protection

The packet storm protection feature allows the user to set limits on the reception rate of broadcast, multicast and destination lookup failure packets. The software allows separate limits to be set for each port, beyond which each of the different packet types are discarded. The software also allows separate limits to be set for each of the packet types. Which of these options can be implemented depends on the model of switch hardware.

By default, packet storm protection is disabled. It can be enabled, and each of the limits can be set using the command:

```
SET SWITCH PORT=port-list [BCLIMIT={NONE | limit}]
[DLFLIMIT={NONE | limit}] [MCLIMIT={NONE | limit}]
```

The BCLIMIT parameter specifies a limit on the rate of reception of broadcast packets for the port(s). The value of this parameter represents a per second rate of packet reception above which packets will be discarded, for broadcast packets. If the value NONE or 0 is specified, then packet rate limiting for broadcast packets is turned off. If any other value is specified, the reception of broadcast packets will be limited to that number of packets per second. See the note below for important information about packet rate limiting. The default value for this parameter is NONE.

The DLFLIMIT parameter specifies a limit on the rate of reception of destination lookup failure packets for the port. The value of this parameter represents a per second rate of packet reception above which packets will be discarded, for destination lookup failure packets. If the value NONE or 0 is specified, then packet rate limiting for destination lookup failure packets is turned off. If any other value is specified, the reception of destination lookup failure packets will be limited to that number of packets per second. See the note below for important information about packet rate limiting. The default value for this parameter is NONE.




---

*A destination lookup failure packet is one for which the switch hardware does not have a record of the destination address of the packet, either Layer 2 or Layer 3 address. These packets are passed to the CPU for further processing, so limiting the rate of reception of these packets may be a desirable feature to improve system performance.*

---

The LEARN parameter specifies whether or not the security feature of limiting the number of MAC addresses learned on these port(s) is enabled. If NONE or 0 is specified, there is no limit set on the number of MAC addresses learned on this port. If a number from 1 to 256 is specified, the switch will stop learning MAC addresses on these port(s) once the number of MAC addresses has been reached. Packets received from other addresses after this time are dealt with as intrusion packets (see the LOCKACTION parameter). The default value for this parameter is NONE.

The MCLIMIT parameter specifies a limit on the rate of reception of multicast packets for the port. The value of this parameter represents a per second rate of packet reception above which packets will be discarded, for multicast packets. If the value NONE or 0 is specified, then packet rate limiting for multicast packets is turned off. If any other value is specified, the reception of multicast packets will be limited to that number of packets per second. See the note below for important information about packet rate limiting. The default value for this parameter is NONE.




---

*The ability of the switch to limit packet reception rates for different classes of packets is dependent on the particular switch hardware. In particular, groups of ports may have to have the same limits set, and the same limit may be set for the different types of packets, depending on the hardware. Whenever packet rate limits are set on switches which have this type of constraint, the latest parameter values entered will supersede earlier values. Also, a message when commands are entered will clearly indicate the effect of the command in cases where parameters for other ports have changed.*

---

On the Rapier switch, packet storm protection limits cannot be set for each individual port, but can be set for each processing block of ports. On the Rapier 24 switch the processing blocks are ports 1-8, 9-16, 17-24, and a processing block each for the uplink ports 25 and 26. The Rapier 24 switch only allows one limit to be set for all three packet types, while allowing each of the packet types to be either limited to this value, or unlimited (NONE).

The SHOW SWITCH PORT command displays the packet storm protection settings (Figure 8 on page 19).

```
SHOW SWITCH PORT=port-list
```

## Port mirroring

Port mirroring allows traffic being received and transmitted on a switch Ethernet port to be sent to another port, usually for the purposes of capturing the data with a protocol analyser. The mirror port to which traffic is sent is set using the command:

```
SET SWITCH MIRROR={NONE|port}
```

The mirror port cannot be part of a trunk group.

Traffic received on a port, traffic transmitted, or both can be mirrored. This is specified, along with the source port(s) from which traffic is sent to the mirror port, using the command:

```
SET SWITCH PORT=port-list MIRROR={NONE|RX|TX|BOTH}
```

The MIRROR parameter specifies the role of these port(s) as a source of mirror traffic. If the value NONE is specified, no traffic received or sent on these port(s) will be mirrored. If the value RX is specified, all traffic received on these port(s) will be mirrored. If the value TX is specified, all traffic transmitted on these port(s) will be mirrored. If the value BOTH is specified, all traffic received and transmitted will be mirrored. Traffic will actually only be mirrored if there is a mirror port defined and the mirror feature is enabled. The default is NONE.

By default mirroring is disabled, no mirror port is set, and no source ports are set to be mirrored. Mirroring is enabled and disabled using the commands:

```
ENABLE SWITCH MIRROR
```

```
DISABLE SWITCH MIRROR
```

## Virtual LANs

A Virtual LAN is a software-defined broadcast domain. The switch's VLAN feature allows the network to be segmented by software management, improving network performance. Workstations, servers, and other network equipment connected to the switch can be grouped according to similar data and security requirements. Several VLANs can be connected to the same switch. Devices that are members of a VLAN only exchange data with each other through the switching capabilities of the switch. Further flexibility can be gained by using VLAN tagging. To exchange data between devices in separate VLANs, the switch's routing capabilities are used.

By default the switch is configured to include all ports in a single default port-based VLAN, with no VLAN tagging required on incoming frames, or added to outgoing frames. If all the devices on the physical LAN are to belong to the same logical LAN, that is, the same broadcast domain, then the default settings will be acceptable, and no additional VLAN configuration is required.

The ability to decouple logical broadcast domains from the physical wiring topology offers several advantages, which include:

- Workstations can be grouped logically or functionally, regardless of their physical location on the network.
- VLAN memberships can be changed at any time by software configuration, without moving the workstations physically, or by simply moving a cable from one port to another.
- By using VLAN tagging, network servers or other network resources can be shared between different work groups without losing data isolation or security.
- One port on the switch can be configured as an uplink to another 802.1Q-compatible switch. By using VLAN tagging, this one port can carry traffic from all VLANs on the switch. (With port based VLANs, one uplink port is required to uplink each VLAN to another switch.)

Two main types of VLAN can be configured on the switch:

- VLANs that are simple logical groupings of ports, that do not use VLAN tags on the frames they receive or send.
- VLANs that add tags to frames transmitted over some ports. A port can belong to more than one tagged VLAN, so that a single port can be used to uplink several VLANs to another compatible switch.

A VLAN can contain a mixture of VLAN tagged and untagged ports.

The switch is VLAN aware, in that it can accept VLAN tagged frames, and supports the VLAN switching required by such tags. A network can contain a mixture of VLAN aware devices, for instance other 802.1Q compatible switches, and VLAN unaware devices, for instance, workstations and legacy switches that do not support VLAN tagging.

The switch can be configured to send VLAN tagged or untagged frames on each port, depending on whether or not the devices connected to the port are VLAN aware. Each port must always belong to at least one VLAN. A port can be untagged for at most one VLAN, and at the same time be tagged for several other VLANs. This means the same port may send both tagged and untagged frames.

A port can belong to only one Spanning Tree entity (STP), and STP membership is per VLAN. A port cannot be added to a VLAN that is in a different STP from the VLANs to which the port already belongs, with one exception. The exception is that an untagged port in the default VLAN that is not tagged for any other VLANs can be moved from the default VLAN to any other VLAN in any STP.

## Creating VLANs without VLAN tags

VLANs that do not send any VLAN-tagged frames are logical groupings of ports. Any devices connected to the member ports share a common broadcast domain. The switch only forwards the traffic in a VLAN to the member ports.

The switch has one default VLAN, which is created when the switch is powered up. More such VLANs can be created on the switch at any time. Each new VLAN is created with a VLAN name that is unique in the switch, and a VLAN Identifier (VID) that uniquely identifies the VLAN on the physical LAN. The default VLAN always has a VID of 1.

VLANs are created and destroyed with the commands:

```
CREATE VLAN=vlannname VID=2..4094  
DESTROY VLAN={vlannname|2..4094|ALL}
```

The VLAN parameter specifies a unique name for the VLAN. This name can be more meaningful than the VID, to make administration easier. The VLAN name is only used within the switch; it is not transmitted to other VLAN-aware devices, or used in the Forwarding Process or kept in the Forwarding Database.

The VID parameter specifies a unique VLAN IDentifier for the VLAN. If VLAN-tagged ports are added to this VLAN, the specified VID is used in the VID field of the tag in outgoing frames. If VLAN-untagged ports are added to this VLAN, the specified VID only acts as an identifier for the VLAN in the Forwarding Database. The default port based VLAN has a VID of 1.

By default, all ports on the switch belong to the default VLAN. Any untagged ports in the default VLAN port can be added untagged to another VLAN, and are then automatically removed from the default VLAN. A port can only be untagged for one VLAN. An untagged port deleted from a VLAN is returned to the default VLAN if it is not a tagged member of any other VLANs.

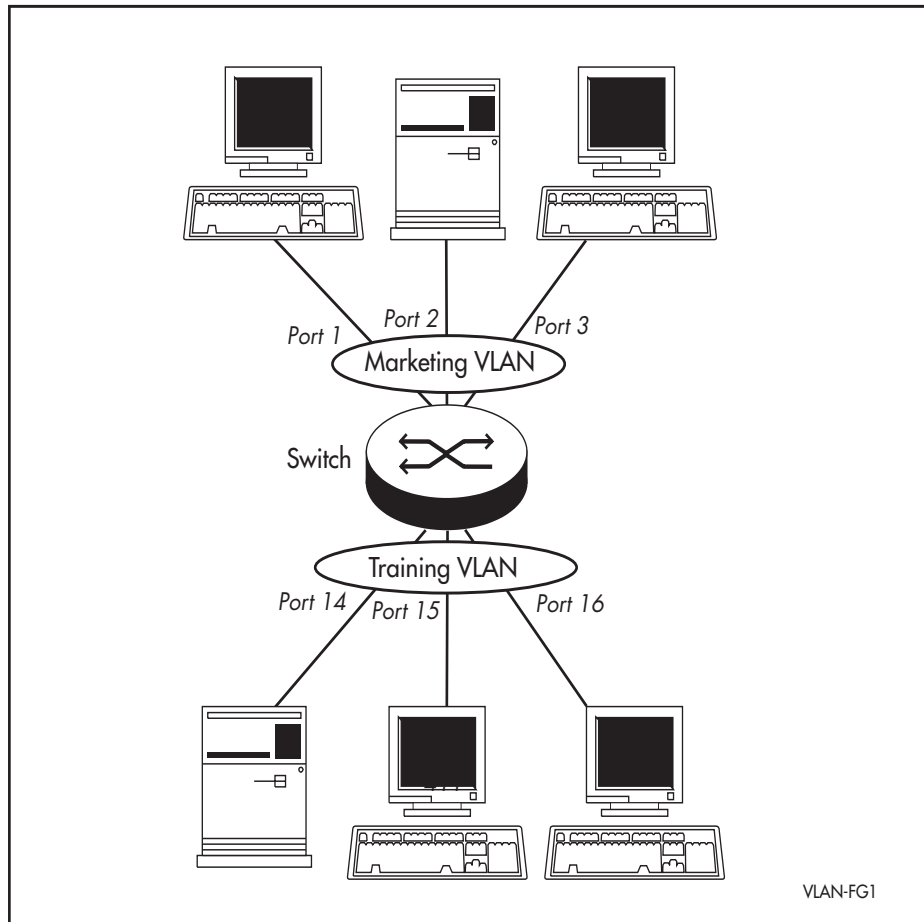
Untagged ports are added to and removed from VLANs with these commands:

```
ADD VLAN={vlannname|2..4094} PORT=port-list  
DELETE VLAN={vlannname|2..4094} PORT=port-list
```

As a VLAN-untagged frame is admitted on a port, the VID of the VLAN for which the port is untagged is associated with the frame. This VID is used to forward the frame only to the ports in the same VLAN.

Figure 10 shows two VLANs. Ports 1-3 belong to one broadcast domain, the Marketing VLAN and ports 14-16 belong to another broadcast domain, the Training VLAN. The switch acts as two separate bridges: one that forwards between the ports belonging to the Marketing VLAN, and a second one that forwards between the ports belonging to the Training VLAN. Devices on in the Marketing VLAN can only communicate with devices in the Training VLAN by using the switch's routing functions.

Figure 10: VLANs with untagged ports



To display the VLANs configured on the switch, use the command:

```
SHOW VLAN [= { vlanname | 1..4094 | ALL }]
```

Figure 11: Example output from the SHOW VLAN command.

```

VLAN Information
-----
Name                : default
Identifier           : 1
Untagged Ports      : 1-10, 13
Tagged Ports        : none
Spanning Tree       : default
Attachments:
Module      Protocol      Format      Discrim      MAC address
-----
IP          IP            Ethernet   0800         -
IP          AR            Ethernet   0806         -
-----

```

Table 8: Parameters displayed in the output of the SHOW VLAN command

Parameter	Meaning
Name	The name of the VLAN.
Identifier	The numerical VLAN Identifier of the VLAN.
Untagged Ports	A list of untagged ports that belong to the VLAN.

**Table 8: Parameters displayed in the output of the SHOW VLAN command**

Parameter	Meaning
Tagged Ports	A list of tagged ports that belong to the VLAN.
Spanning Tree	The name of the Spanning Tree entity to which the VLAN belongs.
<b>Attachments</b>	This section shows information about other modules and protocols using the VLAN module.
Module	The name of the software module attached to the VLAN.
Protocol	The name of the protocol, which is determined from the format and discriminator.
Format	The encapsulation format specified by the module.
Discrim	The discriminator specified by the module to identify which packets of the given format should be received.
MAC Address	The Media Access Control source address for which the module wishes to receive packets. This is commonly known as the Ethernet address.

There are some disadvantages to using VLANs with untagged ports only:

- It is difficult to share network resources, such as servers and printers, across several VLANs. The routing functions in the switch must be configured to interconnect using untagged ports only.
- A VLAN that spans several switches requires a port on each switch for the interconnection of the various parts of the VLAN. If there are several VLANs in the switch that span more than one switch, then many ports are occupied with connecting the VLANs, and so are unavailable for other devices.

These disadvantages can be overcome with the versatility of VLAN tagging.

## VLAN tagging

VLAN tagging provides the advantages of more efficient and flexible use of switch ports and network resources, while maintaining the level of security given by port-based VLANs. With VLAN tagging, a port can belong to several VLANs. This means that network resources can be shared between different VLANs by configuring their ports to belong to more than one VLAN. Only one port is required on each switch to uplink (trunk) all VLAN traffic between two VLAN aware switches, as this port can be configured to belong to all VLANs on the switch.

Support for VLAN tagging is implemented in the switch according to IEEE Standard 802.1Q. Just as with untagged ports, tagged ports in a VLAN belong to the VLAN's broadcast domain. A VLAN Identifier (VID) is defined for each VLAN, and this VID is used to switch traffic through a VLAN aware network so that frames are only transmitted on ports belonging to the VLAN. Other vendors VLAN aware devices on the network can be configured to accept traffic from one or more VLANs. A VLAN-aware server can be configured to accept traffic from many different VLANs, and then return data to each VLAN without mixing or leaking data into the wrong VLANs.

Every frame admitted by the switch has a VID associated with it, either because it already had a VLAN tag when it arrived, or because the VLAN for which the incoming port is untagged was associated with it when it was

admitted. The switch only forwards the frame over those ports that belong to the VLAN specified by this VID. When the switch forwards a frame over a port to another VLAN-aware device (for instance, another switch), it adds a VLAN tag (the same VID) to the frame. When it forwards the frame over a port to a VLAN-unaware device, it transmits it as a VLAN-untagged frame, not including the VID in the frame.

VLANs to be used with VLAN tags are created and destroyed in the same way as VLANs with only untagged ports, with the commands:

```
CREATE VLAN=vlanname VID=2..4094
DESTROY VLAN={vlanname|2..4094|ALL}
```

A tagged VLAN may have VLAN-aware devices connected to some ports that require VLAN tags and legacy devices connected to other ports that cannot accept VLAN tags. Whether VLAN tagged or untagged frames are transmitted on a port to members of a particular VLAN is determined when the ports are added to that VLAN. Ports are added to and deleted from VLANs with these commands:

```
ADD VLAN={vlanname|2..4094} PORT=port-list
[FRAME=TAGGED|UNTAGGED]
DELETE VLAN={vlanname|2..4094} PORT=port-list
```

The FRAME parameter specifies whether VLAN tag headers are included in frames transmitted on the specified ports. If TAGGED is specified, a VLAN tag is added to frames prior to transmission. The port is then called a tagged port for this VLAN. If UNTAGGED is specified, the frame is transmitted without a VLAN tag. The port is then called an untagged port for this VLAN. A port can be untagged for one and only one of the VLANs to which it belongs, or for none of the VLANs to which it belongs. A port can have the FRAME parameter set to TAGGED for zero or more VLANs to which it belongs. It is not possible to add an untagged port to a VLAN if the port is already present in any other port-based VLAN except the default VLAN. If the port is an untagged member of the default VLAN, adding it untagged to another VLAN deletes it from the default VLAN. The default setting is UNTAGGED.

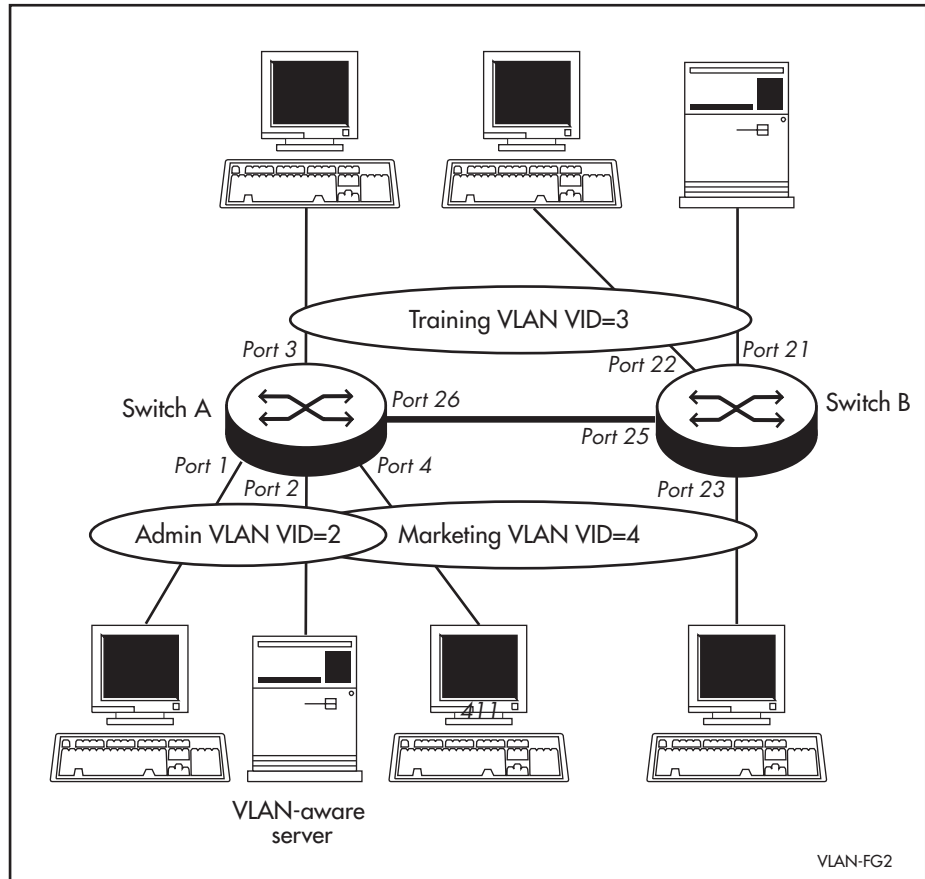
A port can be tagged (send VLAN tagged frames) for some VLANs, and untagged for one VLAN. A VLAN can have both tagged and untagged ports. But a particular port cannot receive and send both tagged and untagged frames for the same VLAN (that is, it cannot be added to the same VLAN as both a tagged and an untagged port).

Figure 12 shows a network that must be configured with VLAN tagging, since the server on port 2 on Switch A belongs to both the Admin and Marketing VLANs. Using VLAN tags, port 25 on Switch A and port 26 belong to both the Marketing VLAN and the Training VLAN, so devices on both VLANs can use this uplink.



*Note that ports tagged for some VLANs and left in the default VLAN as untagged ports will transmit broadcast traffic for the default VLAN. If this is not required, the unnecessary traffic in the switch can be reduced by actively creating another VLAN on the switch and adding the untagged ports that are intended to be active to it, so that they no longer generate broadcast traffic to the default VLAN.*

Figure 12: Tagged VLANs



To display the VLANs configured on the switch, use the command:

```
SHOW VLAN [= { vlanname | 1..4094 | ALL }]
```

Some additional information can be displayed which may help with trouble shooting your network. VLAN debugging mode is disabled by default, and can be enabled for a specified time, disabled, and displayed using the commands:

```
ENABLE VLAN [= { vlanname | 1..4094 }] DEBUG = { PKT | ALL }
[TIMEOUT = { 1..4000000000 | NONE }] [OUTPUT = CONSOLE]

DISABLE VLAN [= { vlanname | 1..4094 }] DEBUG = { PKT | ALL }

SHOW VLAN DEBUG
```

The OUTPUT parameter set to CONSOLE specifies that the debugging information produced is sent to the console. Use this parameter if the ENABLE VLAN DEBUG command is used in a script. The debugging data is by default sent to the port on which it received the ENABLE VLAN DEBUG command. Use this option if the command is used in a script, since a script is not received on a port.

The TIMEOUT parameter specifies the time in seconds for which any debugging will be enabled on the specified VLAN. This reduces the risk of the switch and the display being overloaded with too much debugging information. This value overrides any previous VLAN debugging timeout values for VLAN debugging, even if they were specified for other debugging modes. If TIMEOUT is not specified, the time out is the most recent TIMEOUT value used set in an ENABLE VLAN DEBUG command, or NONE if it has not been previously set.

**Figure 13: Example output from the SHOW VLAN DEBUG command**

Vlan	Enabled Debug Modes	Output	Timeout
Vlan1	PKT	16	60
Vlan	Enabled Debug Modes	Output	Timeout
Vlan4094	None		

**Table 9: Parameters in the output of the SHOW VLAN DEBUG command**

Parameter	Meaning
VLAN	The VLAN Identifier of the VLAN, shown after the constant "VLAN".
Enabled Debug Modes	The debugging option for the VLAN; one of "PKT" or "NONE".
Output	The output device for the VLAN.
Timeout	The value of the timeout in seconds.

## Summary of VLAN tagging rules

When designing a VLAN and adding ports to VLANs, the following rules apply.

1. A port must belong to at least one VLAN.
2. A port can be untagged for zero or one VLAN. A port that is untagged for a VLAN transmits frames destined for that VLAN without a VLAN tag in the Ethernet frame.
3. A port can be tagged for zero or more VLANs. A port that is tagged for a VLAN transmits frames destined for that VLAN with a VLAN tag, including the VID of the VLAN.
4. A port cannot be untagged and tagged for the same VLAN.

## The Switching Process

The switching process comprises related but separate processes. The *Ingress Rules* admit or discard frames based on their VLAN tagging. The *Learning Process* learns the MAC addresses and VLAN membership of frames admitted on each port. The *Forwarding Process* determines which ports the frames are forwarded to, and the *Quality of Service* priority with which they are transmitted. Finally, the *Egress Rules* determine for each frame whether VLAN tags are included in the Ethernet frames that are transmitted. These processes assume that each station on the extended LAN has a unique data link layer address, and that all data link layer frames have a header which includes the source MAC address and destination MAC address.

## The Ingress Rules

When a frame first arrives at a port, the *Ingress Rules* for the port check the VLAN tagging in the frame to determine whether it will be discarded or forwarded to the Learning Process.

The first check depends on whether the *Acceptable Frames* parameter is set to *Admit All Frames* or to *Admit Only VLAN Tagged Frames*. If it is set to *Admit Only VLAN Tagged Frames*, then any incoming frames with a null VLAN Identifier (VID) are discarded. If a port is tagged for every VLAN it belongs to, then it is automatically set to *Admit Only VLAN Tagged Frames*. Frames with a null VID are VLAN-untagged frames, or frames with priority tagging only.

The switch requires a VLAN Identifier (VID) for every frame it switches. If all frames are admitted by the *Acceptable Frames* parameter, the second part of the *Ingress Rules* associates with each untagged frame admitted the VID of the VLAN for which the port is untagged.

Every port belongs to one or more VLANs, and now every incoming frame has a VID to show which VLAN it belongs to. The final part of the *Ingress Rules* depends on whether *Ingress Filtering* is enabled for the port. If *Ingress Filtering* is disabled, all frames are passed on to the Learning Process, regardless of which VLAN they belong to. If *Ingress Filtering* is enabled, frames are admitted only if they have the VID of a VLAN to which the port belongs. If they have the VID of a VLAN to which the port does not belong, they are discarded.

The default settings for the switch are for all ports to be untagged members of the default VLAN, and for the *Ingress Rules* are to *Admit All Frames* on ports that are untagged for one VLAN. *Ingress Filtering* is OFF by default. This means that if no VLAN configuration has been done, all incoming frames pass on to the Learning Process, regardless of whether or not they are VLAN tagged. If a port is tagged for every VLAN to which it belongs, only VLAN tagged frames are admitted, and this setting cannot be changed. The parameters for each port's *Ingress Rules* can be configured using the command:

```
SET SWITCH PORT=port-list [ACCEPTABLE={VLAN|ALL}]  
[INFILTERING={ON|OFF}] [other-parameters...]
```

The ACCEPTABLE parameter sets the *Acceptable Frames Type* parameter, in the *Ingress Rules*, which controls reception of VLAN-tagged and VLAN-untagged frames on the port. If ALL is specified, then the *Acceptable Frames Type* parameter is set to *Admit All Frames*. If VLAN is specified, the parameter is set to *Admit Only VLAN-tagged Frames*, and any frame received that carries a null VLAN Identifier (VID) is discarded by the ingress rules. Untagged frames and priority-tagged frames carry a null VID. VLAN-untagged frames admitted according to the ACCEPTABLE parameter have associated with them the VID of the VLAN for which they are an untagged port. The ACCEPTABLE parameter can only be set if the port is untagged for one VLAN. In this case, the default is ALL, admitting all tagged and untagged frames. If the port is tagged for all the VLANs to which it belongs, the ACCEPTABLE parameter is automatically set to VLAN, and cannot be changed to admit untagged frames.

The INFILTERING parameter enables or disables *Ingress Filtering* on the specified ports of frames admitted according to the ACCEPTABLE parameter. Each port on the switch belongs to one or more VLANs. If INFILTERING is set to ON, *Ingress Filtering* is enabled: any frame received on a specified port is only admitted if the port belongs to the VLAN with which the frame is associated. Conversely, any frame received on the port is discarded if the port does not belong to the VLAN with which the frame is associated. VLAN-

untagged frames admitted by the ACCEPTABLE parameter are admitted, since they have the VID of the VLAN for which the port is an untagged member. If OFF is specified, Ingress Filtering is disabled, and no frames are discarded by this part of the Ingress Rules. The default setting is OFF.

To display the current Ingress Rules (Figure 8 on page 19), use the command:

```
SHOW SWITCH PORT=port-list
```

## The Learning Process

The Learning Process uses an *adaptive learning* algorithm, sometimes called *backward learning* to discover the location of each station on the extended LAN.

All frames admitted by the Ingress Rules on any port are passed on to the Forwarding Process if they are for destinations within the same VLAN. Frames destined for other VLANs are passed to the Layer 3 protocol, for instance IP. For every frame admitted, the frame's source MAC address and VID are compared with entries in a Forwarding Database for the VLAN (also known as a MAC address table, or a forwarding table) maintained by the switch. The Forwarding Database contains one entry for every unique station MAC address the switch knows in each VLAN.

If the frame's source address is not already in the Forwarding Database for the VLAN, the address is added and an ageing timer for that entry is started. If the frame's source address is already in the Forwarding Database, the ageing timer for that entry is restarted. By default, switch learning is enabled, and it can be disabled or enabled using the commands:

```
DISABLE SWITCH LEARNING
```

```
ENABLE SWITCH LEARNING
```

If the ageing timer for an entry in the Forwarding Database expires before another frame with the same source address is received, the entry is removed from the Forwarding Database. This prevents the Forwarding Database from being filled up with information about stations that are inactive or have been disconnected from the network, while ensuring that entries for active stations are kept alive in the Forwarding Database. By default, the ageing timer is enabled, and it can be disabled or enabled using the commands:

```
ENABLE SWITCH AGEINGTIMER
```

```
DISABLE SWITCH AGEINGTIMER
```

The default value of the ageing timer is 300 seconds (5 minutes), and this can be modified using the command:

```
SET SWITCH AGEINGTIMER=10..1000000
```

The Forwarding Database relates a station's (source) address to a port on the switch, and is used by the switch to determine from which port (if any) to transmit frames with a destination MAC address matching the entry in the station map.

## The Forwarding Process

The Forwarding Process forwards received frames that are to be relayed to other ports in the same VLAN, filtering out frames on the basis of information contained in the station map and on the state of the ports. If a frame is received on the port for a destination in a different VLAN, it is either Layer 3 switched if it is an IP packet, or looked up in the Layer 3 routing tables (see the *AR Series Router Reference Manual* at <http://www.alliedtelesyn.co.nz/support/rapier/>.)

Forwarding occurs only if the port on which the frame was received is in the Spanning Tree 'Forwarding' state. The destination address is then looked up in the Forwarding Database for the VLAN. If the destination address is not found, the switch floods the frame on all ports in the VLAN except the port on which the frame was received. If the destination address is found, the switch discards the frame if the port is not in the STP 'Forwarding' state, if the destination address is on the same port as the source address, or if there is a static filter entry for the destination address set to DISCARD ("*Filtering*" on page 35). Otherwise, the frame is forwarded on the indicated port.

This whole process can further be modified by the action of static switch filters. These are configurable filters which allow switched frames to be checked against a number of entries.

The Forwarding Process provides storage for queued frames to be transmitted over a particular port or ports. Which transmission queue a frame is sent to is determined by the user priority tag in the Ethernet frame, and the Quality of Service mapping.

## Filtering

The switch has a Forwarding Database stored in RAM, entries in which determine whether frames are forwarded or discarded over each port. Entries in this Forwarding Database are created dynamically by the Learning Process, and by STP (see "*Spanning Tree Protocol (STP)*" on page 38) if these are enabled. Static entries can be configured by the user via the command line. Filtering is specified in the IEEE 802.1D Standard "*Media Access Control (MAC) Bridges*".

Configurable switch filters can be added to each switch port. Switch filters consist of a number of static entries, each of which contains a match condition and a port list. When a frame is received on a port which has a filter configured to it, the frame is checked against all entries in the filter. If a match is found, the port list is used to modify the ports to which the frame can be forwarded, subject to the dynamic filtering entries and the state of the port. The port list can indicate that the frame can be forwarded to all ports in a particular VLAN, to no ports (in which case the frame is immediately discarded) or to a subset of the configured ports in the VLAN.

The Forwarding Database supports queries by the Forwarding Process about whether frames with given values of the destination MAC address field should be forwarded to a given port.

To add or delete a static switch filter entry, use the command:

```
ADD SWITCH FILTER DESTADDRESS=macadd ACTION={FORWARD|DISCARD}
    PORT[=port-list] [ENTRY=entry] [VLAN={vlanname|1..4094}]

DELETE SWITCH FILTER ENTRY=entry
```

To display current switch filter entries, use the command:

```
SHOW SWITCH FILTER [DESTADDRESS=macadd] [ENTRY=entry]
    [PORT=port-list] [VLAN={vlanname|1..4094}]
```

Figure 14: Example output from the SHOW SWITCH FILTER command.

Switch Filters						
Dest.	VLAN	Age	Action	St.	Port	List
00-00-b4-85-97-98	default(0001)	000000	dyn	s	2	
00-00-b4-87-a9-60	default(0001)	000000	dyn	s	1	
00-00-b4-a4-97-7e	accounting(0002)	000000	for	s	1,2-5	
00-00-cd-00-10-81	marketing(0003)	000239	dis	l	7,8,10	
00-e0-29-25-46-7a	default(0001)	000000	dyn	l	1	
.	.	.	.	.	.	.
.	.	.	.	.	.	.
00-f0-c0-2b-69-0b	abcdefghijklmno(4094)	000000	dyn	s	PPP0	

Table 10: Parameters in the output of the SHOW SWITCH FILTER command

Parameter	Meaning
Dest.	The destination MAC address for the entry.
VLAN	The VLAN name and identifier for the entry.
Age	The age in seconds of the filter entry. Static filter entries have the value 000000.
Action	The action specified by the filter entry: one of "for" (forward), "dis" (discard) or "dyn" (dynamic).
St.	The status of the filter entry: one of "s" (static) or "l" (learned).
Port list	The list of outbound ports to match for the filter entry to be applied.

The destination MAC address of a frame to be forwarded is checked against the Forwarding Database. If there is no entry for the destination address the frame is transmitted on all ports, except the port on which the frame was received, which are in the 'Forwarding' state. This process is referred to as *flooding*. If an entry is found in the Forwarding Database, but the entry is not marked as 'Forwarding' or the entry points to the same port the frame was received on, the frame is discarded. Otherwise, the frame is transmitted on the port specified by the entry in the Forwarding Database.

A dynamic entry is automatically deleted from the Forwarding Database when its ageing timer expires.

## Quality of Service

The switch hardware has a number of Quality of Service (QOS) *egress queues* that can be used to give priority to the transmission of some frames over other frames on the basis of their user priority tagging. The user priority field in an incoming frame (with value 0 to 7) determines which of the eight priority levels the frame is allocated. When a frame is forwarded, it is sent to a QOS egress queue on the port determined by the mapping of priority levels to QOS egress queues. All frames in the first QOS queue are sent before any frames in the second QOS egress queue, and so on, until frames in the last QOS egress queue,

which are only sent when there are no frames waiting to be sent in any of the higher QOS egress queues.

Which traffic class is sent to which QOS egress queue can be configured with this command:

```
SET SWITCH QOS=P1 , P2 , P3 , P4 , P5 , P6 , P7 , P8
```

The Rapier 24 has four QOS egress queues. It has a default mapping of priority levels to QOS egress queues as defined in *IEEE Standard 802.1Q* (Table 11).

**Table 11: Default priority level to queue mapping for four QOS egress queues**

Priority level	QOS Egress Queue
0	1
1	0
2	0
3	1
4	2
5	2
6	3
7	3

To display the mapping of user priority to QOS egress queues, use the command:

```
SHOW SWITCH QOS
```

**Figure 15: Example output from the SHOW SWITCH QOS command**

```

Priority Level QOS egress queue
=====
0 ..... 1
1 ..... 0
2 ..... 0
3 ..... 1
4 ..... 2
5 ..... 2
6 ..... 3
7 ..... 3

```

**Table 12: Parameters in the output of the SHOW SWITCH QOS command**

Parameter	Meaning
Priority level	The priority level of the frame.
QOS egress queue	The Quality Of Service egress queue that frames with this priority level join.

## The Egress Rules

Once the Forwarding Process has determined which ports and transmission queues to forward a frame from, the Egress Rules for each port determine whether or not the outgoing frame is VLAN-tagged with its VID. (See “*Virtual LANs*” on page 25).

Whether outgoing frames for a VLAN are tagged when transmitted from the port is configured when the port is added to the VLAN, and can be changed later, using this commands:

```
ADD VLAN={vlanname|1..4094} PORT={port-list|ALL}
    [FRAME={TAGGED|UNTAGGED}]

SET VLAN={vlanname|1..4094} PORT={port-list|ALL}
    FRAME={UNTAGGED|TAGGED}
```

## Spanning Tree Protocol (STP)

The Spanning Tree Protocol (STP) makes it possible to automatically disable redundant paths in a network to avoid loops, and enable them when a fault in the network means they are needed to keep traffic flowing. A sequence of LANs and switches may be connected together in an arbitrary physical topology resulting in more than one path between any two switches. If a loop exists, frames transmitted onto the extended LAN would circulate around the loop indefinitely, decreasing the performance of the extended LAN. On the other hand, multiple paths through the extended LAN provide the opportunity for redundancy and backup in the event of a bridge experiencing a fatal error condition.

The spanning tree algorithm ensures that the extended LAN contains no loops and that all LANs are connected by:

- Detecting the presence of loops and automatically computing a logical loop-free portion of the topology, called a *spanning tree*. The topology is dynamically pruned to a spanning tree by declaring a port redundant, and placing the port into a non-‘Forwarding’ state.
- Automatically recovering from a switch failure that would partition the extended LAN by reconfiguring the spanning tree to use redundant switches.

The logical tree computed by the spanning tree algorithm has the following properties:

- A single switch, called the *root bridge*, forms a unique root to the tree. The root bridge is the bridge with the lowest Bridge ID. Each switch in an extended LAN is uniquely identified by its Bridge ID, which comprises the switch’s root priority (a spanning tree parameter) and its MAC address.
- Each switch or LAN in the tree, except the root bridge, has a unique parent.
- The unique parent of a LAN is the *designated bridge* for the LAN. Each LAN has a single switch, called the designated bridge, that logically connects the LAN to the next LAN closer to the root bridge. Each port connecting a switch to a LAN has an associated *cost*. The *root path cost* is the sum of the costs for each port between the switch and the root bridge. The designated bridge for a LAN is the switch on the LAN with the lowest root path cost, and therefore logically closer to the root bridge. If two switch’s on the same LAN have the same lowest root path cost, the switch with the lowest bridge ID is elected the designated bridge.

- The unique parent of a switch is the LAN, to which the switch is attached, that is closest to the root bridge.

The spanning tree computation is a continuous, distributed process. The algorithm uses the following steps to establish the spanning tree:

1. A unique *root* bridge is elected by the switches in the LAN.
2. A designated bridge is elected for each LAN in the extended LAN by the switches in the LAN.
3. The logical spanning tree is computed and redundant paths are removed.

Once the spanning tree is established, it is maintained by:

1. Replacing a failed port with a redundant backup port.
2. Detecting and removing loops by declaring a port redundant and removing it from the logical spanning tree.
3. Maintaining address timers that control the ageing of station map address entries.

The logical spanning tree, sometimes called the active topology, includes the root bridge and all designated bridges, i.e. all the ports that are to be used for communication within the STP. These ports are dynamically marked as 'Forwarding'. Ports removed from the logical spanning tree are not in the 'Forwarding' state. To implement the spanning tree algorithm, switches communicate with one another using the Spanning Tree Protocol. The primary protocol data unit (PDU) is the *Hello message* or *Configuration Bridge Protocol Data Unit* (BPDU), which includes the following information:

- The bridge ID of the root bridge.
- The distance (or cost) from this switch to the root bridge.
- The bridge ID of the designated bridge on this LAN.

Hello messages are initiated at regular intervals by the root bridge and propagate through the extended LAN.

## Electing the Root Bridge and Designated Bridge

Each spanning tree has a *root bridge*, which initiates the propagation of Hello messages through the extended LAN, and sets the values of parameters that control the spanning tree computation process. The root bridge is the switch with the lowest bridge ID and is elected by the exchange of Hello packets. When a switch receives a Hello packet it compares the value of the root bridge ID in the message to the value of the root bridge ID parameter in its own spanning tree database. If the value in the message is better, the switch stores the new value in its database and sends Hello messages with the new value out on its other ports. Otherwise, the switch continues to send out Hello messages with the value currently stored in its spanning tree database. By this process all switches in the extended LAN will eventually learn the bridge ID of the root bridge.

Each LAN has a single switch, called the *designated bridge*, that logically connects the LAN to the next LAN closer to the root bridge. The designated bridge for a LAN is the switch on the LAN with the lowest root path cost and bridge ID. The designated bridge is elected by the exchange of Hello messages, in the same way that the root bridge is elected. The election of a new root bridge, or a switch becoming unavailable due to a fatal error condition, will normally result in the election of a new designated bridge in the next few rounds of Hello messages.

## Switch port states

A switch port may be in one of five states (Table 13), determined dynamically by STP (“*Spanning Tree Protocol (STP)*” on page 38).

**Table 13: Switch port states**

State	Meaning
DISABLED	Switching operations are disabled on the port. In particular, the Forwarding Process and the Spanning Tree entity are disabled for transmit and receive operations on the port.
LISTENING	The port is enabled for receiving frames only.
LEARNING	The port is enabled for receiving frames only, and the Forwarding Process is placing new source address information in the station map.
FORWARDING	The normal state for a switch port. The Forwarding Process and the Spanning Tree entity are enabled for transmit and receive operations on the port.
BLOCKING	The Spanning Tree entity has disabled the Forwarding process for transmit and receive operations on the port, but the Spanning Tree entity itself remains enabled for transmit and receive operations on the port.

To display the state of the about switch ports (Figure 17 on page 45), use the command:

```
SHOW STP PORT=port-list
```

## Multiple Spanning Trees

A Rapier switch in default LAN configuration has a *default* Spanning Tree enabled, spanning only a single default VLAN, to which all ports belong. The switches in the LAN run a distributed Spanning Tree Algorithm to create the Spanning Tree.

In a network with more VLANs configured, all VLANs belong by default to the default Spanning Tree called *default*. This configuration will suit some LANs. In order to distribute traffic more evenly over multiple paths, Multiple Spanning Trees can be created, with each Spanning Tree encompassing multiple VLANs. Spanning Tree Protocol entities, called STPs here, operate independently of each other; each STP has its own Root Bridge and active path. Once an STP is created, one or more VLANs can be assigned to it.

If creating multiple STPs in a network, consider the following:

- A new instance of the Spanning Tree Protocol (STP) need not be created identical to the topology of the VLAN(s). All VLANs are aligned along the Spanning Tree from which they are formed; a given VLAN is defined by a subset of the topology of the Spanning Tree upon which it operates.
- The topology of the VLAN is dynamic. The structure of the VLAN may change due to new devices requesting or releasing the services available via the VLAN.
- Any VLAN can only belong to a single STP.
- Any port in the network must only belong to a single STP. If a port is a member of multiple VLANs, then all those VLANs must belong to the

same STP. Within any given STP, all VLANs belonging to it use the same Spanning Tree.

- The topology of the VLAN is dynamic. The structure of the VLAN may change due to new devices requesting or releasing the services available via the VLAN. The dynamic nature of VLANs has the advantages of flexibility and bandwidth conservation, at the cost of network management complexity.



*Multiple Spanning Trees in a VLAN environment have not been standardised. This means the BPDUs used in communicating Spanning Tree Algorithm parameters are vendor specific, since the BPDUs defined by IEEE std 802.1D are based on a single Spanning Tree and do not have any reference to VLANs. A LAN comprising switches from more than one vendor should, therefore, only use the single default STP. Multiple STPs should not be created.*

## Configuring STP

By default, the switch has one *default* STP which cannot be destroyed. In most situations this default STP will suffice. However, further instances of the Spanning Tree Protocol (STPs) can be created and destroyed using the commands:

```
CREATE STP=stpname
DESTROY STP={stpname|ALL}
```

By default, all VLANs, and therefore all ports, belong to the *default* STP. To add or delete a VLAN and all the ports belonging to it from any other STP, use the commands:

```
ADD STP=stpname VLAN={vlanname|2..4094}
DELETE STP=stpname VLAN={vlanname|2..4094|ALL}
```

The default STP is enabled by default at switch start up, while STPs created by a user are disabled by default when they are created. To enable or disable STPs, use the commands:

```
ENABLE STP[=stpname|ALL]
DISABLE STP[=stpname|ALL]
```

The Spanning Tree Protocol uses three configurable parameters for the time intervals that control the flow of the STP information on which the dynamic STP topology depends: the HELLOTIME, FORWARDDELAY and MAXAGE parameters. All switches in the STP use the values for these parameters sent by the root bridge and forwarded by the designated bridges. A particular switch may be configured with different time intervals, which would be used by the whole spanning tree if this switch became the root bridge.

The HELLOTIME parameter, with a default value of 2 seconds, determines how often the switch sends Hello messages containing spanning tree configuration information if it is the *root bridge*, or is trying to become the root bridge in the network. Setting a shorter value for HELLOTIME than the default of 2 seconds makes the network more robust; setting a longer time uses less processing overheads.

The MAXAGE parameter, with a default of 20 seconds, determines the maximum time that dynamic STP configuration information is stored in the switch, before it is considered too old, and discarded. The value can be set at approximately two seconds for every hop across the network. If this value is too small, the STP may sometimes configure unnecessarily. If it is too long,

there can be delays in adapting to a change in the topology, for instance when a fault occurs.

The FORWARDDELAY parameter is used to prevent temporary loops in the network occurring in the briefly unstable topology while a topology change is propagated through the network. When a port that has been in the Blocking state in a particular STP topology is to move into the Forwarding state after a topology change, it must first pass through the Listening and Learning states, during which it cannot receive or transmit packets. The FORWARDDELAY parameter determines how long the ports remains in each of the Listening and Learning states before moving on to the Forwarding state in the active topology, that is half the time between when it is decided that the port will become part of the spanning tree, and when it is allowed to forward traffic. The FORWARDDELAY parameter should be at least half the time it takes for a topology change message to reach the whole network. A value that is too short risks the temporary creation of loops, which can seriously degrade switch performance. A longer value can result in delays in the network after topology changes. The default FORWARDDELAY value is 15 seconds.

To modify the parameters controlling these time intervals, use the command:

```
SET STP={stpname|ALL} [FORWARDDELAY=4..30] [HELLOTIME=1..10]
[MAXAGE=6..40] [other-parameters...]
```

The value of the PRIORITY parameter is used to set the writable portion of the bridge ID, i.e. the first two octets of the (8-octet long) Bridge Identifier. The remaining 6 octets of the bridge ID are given by the MAC address of the switches. The Bridge Identifier parameter is used in all configuration Spanning Tree Protocol packets transmitted by the switch. The first two octets, specified by the PRIORITY parameter, determine the switches priority for becoming the *root bridge* or a *designated bridge* in the network, with the lowest number indicating the highest priority. In fairly simple networks, for instance those with a small number of switches in a meshed topology, it may make little difference which switch is selected to be the root bridge, and no modifications may be needed to the default PRIORITY parameter, which has a default value of 32768. In more complex networks, one or more switches are likely to be more suitable candidates for the root bridge role, for instance by virtue of being more central in the physical topology of the network. In these cases the STP PRIORITY parameters for at least one of the switches should be modified.

To change the STP priority value, use the command:

```
SET STP={stpname|ALL} PRIORITY=0..65535 [other-parameters...]
```

The PRIORITY parameter sets the priority of the switch to become the Root Bridge. The lower the value of the Bridge Identifier, the higher the priority. If the PRIORITY parameter is set, either by specifying the PRIORITY or DEFAULT parameters, the specified STP is initialised. Counters for the STP are not affected. The default value for PRIORITY is 32768.

To restore STP timer and priority defaults, use the command:

```
SET STP={stpname|ALL} [DEFAULTS]
```

Changing the STP PRIORITY using either of the previous commands initialises the STP, so that elections for the root bridge and designated bridges begin again, without resetting STP counters. To display general information about STPs on the switch, use the command:

```
SHOW STP [= {stpname|ALL}]
```

**Figure 16: Example output from the SHOW STP command.**

```

Spanning Tree Protocol
-----

Name ..... stp1
VLAN members ..... vlan1 (5)
Status ..... OFF
Number of Ports ..... 2
    Number Enabled ..... 0
    Number Disabled ..... 2
Bridge Identifier ..... 32768 : 00-00-cd-00-a9-a5
Designated Root ..... 32768 : 00-00-cd-00-a9-a5
Max Age ..... 20
Hello Time ..... 2
Forward Delay ..... 15
Bridge Max Age ..... 20
Bridge Hello Time ..... 2
Bridge Forward Delay .. 15
Hold Time ..... 1

Name ..... default
VLAN members ..... default (1)
Status ..... ON
Number of Ports ..... 21
    Number Enabled ..... 0
    Number Disabled ..... 21
Bridge Identifier ..... 32768 : 00-00-cd-00-a9-a5
Designated Root ..... 32768 : 00-00-cd-00-a9-a5
Max Age ..... 20
Hello Time ..... 2
Forward Delay ..... 15
Bridge Max Age ..... 20
Bridge Hello Time ..... 2
Bridge Forward Delay .. 15
Hold Time ..... 1

```

**Table 14: Parameters in the output of the SHOW STP command**

Parameter	Meaning
STP Name	The name of the Spanning Tree.
VLAN members	A list of the VLANs that are members of the STP. VLAN Identifiers are shown in brackets.
Status	The status of the STP; either ON or OFF.
Number of Ports	The number of ports belonging to the STP.
Number Enabled	The number of ports that have been enabled and are being considered by the Spanning Tree Algorithm.
Number Disabled	The number of ports that have been disabled and are not being considered by the Spanning Tree Algorithm.
Bridge Identifier	The unique Bridge Identifier of the switch. This parameter consists of two parts, one of which is derived from the unique Switch Address, and the other of which is the priority of the switch.
Designated Root	The unique Bridge Identifier of the bridge assumed to be the Root.

**Table 14: Parameters in the output of the SHOW STP command**

Parameter	Meaning
Root Port	The port number of the root port for the switch. If the switch is the Root Bridge this parameter is not valid and is not shown.
Root Path Cost	The cost of the path to the Root from this switch. If the switch is the Root Bridge this parameter is not valid and is not shown.
Max Age	The maximum age of received Configuration Message information before it is discarded.
Hello Time	The time interval between successive transmissions of the Configuration Message information by a switch which is attempting to become the Root or which is the Root.
Forward Delay	The time ports spend in the Listening state and Learning state before moving to the Learning or Forwarding state respectively. Also the value used for the ageing timer for the dynamic entries in the Forwarding Database while received Configuration Messages indicate a topology change.
Switch Max Age	The value of the Max Age parameter when the switch is the Root or is attempting to become the Root. This parameter is set by the MAXAGE parameter in the SET STP command.
Switch Hello Time	The value of the Hello Time parameter when the switch is the Root or is attempting to become the Root. This parameter is set by the HELLOTIME parameter in the SET STP command.
Switch Forward Delay	The value of the Forward Delay parameter when the switch is the Root or is attempting to become the Root. This parameter is set by the FORWARDDELAY parameter in the SET STP command.
Hold Time	The minimum time in seconds between the transmission of configuration BPDUs through a given LAN Port. The value of this fixed parameter is 1, as specified in IEEE Std 802.1D.

Each port has a port priority, with a default value of 128, used to determine which port should be the root port for the STP if two ports are connected in a loop. The lowest number indicates the highest priority.

```
SET STP PORT={port-list|ALL} PORTPRIORITY=0..255 [other-parameters...]
```

The PORTPRIORITY parameter sets the value of the priority field contained in the port identifier. The Spanning Tree Algorithm uses the port priority when determining the root port for each switch. The port with the lowest value is considered to have the highest priority. The default value is 128.

Each port also has a path cost, which is used if the port is the root port for the STP on the switch. The path cost is added to the root path cost field in configuration messages received on the port to determine the total cost of the path to the root bridge. The default PATHCOST values depend on the port speed, according to the formula:

$$\text{PATHCOST} = 1000 / \text{Port\_Speed\_in\_MB\_per\_second}$$

so that a port operating at 10Mbps has a default pathcost of 100, a port operating at 100 Mbps has a default pathcost of 10, and a port operating at 1 Gbps has a default pathcost of 1. Setting the pathcost to a larger value on a

particular port is likely to reduce the traffic over the LAN connected to it. This may be appropriate if the LAN has lower bandwidth, or if there are reasons for limiting the traffic across it. To modify the STP port pathcost, use the command:

```
SET STP PORT={port-list|ALL} PATHCOST=1..1000000 [other-parameters...]
```

The PATHCOST parameter sets the path cost for each port. The PATHCOST for a LAN port should be set to a maximum of 65535. If the port is to be the root port then this value is used to determine the total cost from the switch to the Root Bridge. The pathcost can be calculated using the formula:

$$\text{PATHCOST} = 1000 / \text{Port\_Speed\_in\_MB\_per\_second}$$

The default value for PATHCOST is set according to the speed. For a port operating at 100 Mbps, the default value is 19. For a port operating at 10 Mbps, the default value is 100.

To restore default port pathcost and priority, use the command:

```
SET STP PORT={port-list|ALL} DEFAULTS
```

When an STP is enabled in a looped or meshed network, it disables and enables particular ports belonging to it dynamically, to eliminate redundant links. All ports in a VLAN belong to the same STP, and their participation in STP configuration, and hence the possibility of them being elected to the STP's active topology is enabled by default. To enable or disable particular ports, use the commands:

```
ENABLE STP PORT={port-list|ALL}
```

```
DISABLE STP PORT={port-list|ALL}
```

To display STP port information, use the command:

```
SHOW STP PORT[={port-list|ALL}]
```

**Figure 17: Example output from the SHOW STP PORT command.**

```
Spanning Tree Protocol
-----
Port ..... 1
  State ..... Forwarding
  STP ..... default
  STP Status ..... ON
  Pathcost ..... 19
  Designated Cost ..... 0
  Designated Bridge ..... 32768 : 00-00-cd-00-a9-a5
  Port Priority ..... 128

Port ..... 2
  State ..... Disabled
  STP ..... all1
  STP Status ..... OFF
  Pathcost ..... 19
  Designated Cost ..... 0
  Designated Bridge ..... 32768 : 00-00-cd-00-a9-a5
  Port Priority ..... 128
-----
```

**Table 15: Parameters displayed in the output of the SHOW STP PORT command**

Parameter	Meaning
Port	The number of the port.
State	The state of the port; one of "Disabled", "Blocking", "Listening", "Learning" or "Forwarding".
STP	The name of the STP that the port is a member of.
STP Status	The status of the STP that the port is a member of; one of "ON" or "OFF".
Pathcost	The pathcost of the port.
Designated Cost	The designated cost for the port.
Designated Bridge	The unique Bridge Identifier of the switch assumed to be the Root.
Port Priority	The priority of the port.

The spanning tree algorithm can be recalculated at any time, and all timers and counters be initialised, using the command:

```
RESET STP={stpname|ALL}
```

To show STP counters, use the command:

```
SHOW STP[={stpname|ALL}] COUNTER
```

**Figure 17-16: Example output from the SHOW STP COUNTER command**

STP Name: stp3	
Receive:	Transmit:
Total STP Packets ..... 0	Total STP Packets ..... 0
Configuration BPDU .... 0	Configuration BPDU .... 2618
TCN BPDU ..... 0	TCN BPDU ..... 0
Invalid BPDU ..... 0	
Discarded:	
Port Disabled ..... 0	
Invalid Protocol ..... 0	
Invalid Type ..... 0	
Config BPDU length .... 0	
TCN BPDU length ..... 0	

**Table 17: Parameters in the output of the SHOW STP COUNTER command**

Parameter	Meaning
Receive: Total STP Packets	The total number of STP packets received. Valid STP packets comprise Configuration BPDUs and Topology Change Notification (TCN) BPDUs.
Receive: Configuration BPDU	The number of valid Configuration BPDUs received.
Receive: TCN BPDU	The number of valid Topology Change Notification BPDUs received.
Receive:Invalid BPDU	The number of invalid STP packets received.
Transmit: Configuration BPDU	The number of Configuration BPDUs transmitted.

**Table 17: Parameters in the output of the SHOW STP COUNTER command**

Parameter	Meaning
Transmit: TCN BPDU	The number of Topology Change Notification BPDUs transmitted.
Discarded: Port Disabled	The number of BPDUs discarded because the port that the BPDU was received on was disabled.
Discarded: Invalid Protocol	The number of STP packets that had an invalid Protocol Identifier field or invalid Protocol Version Identifier field.
Discarded: Invalid Type	The number of STP packets that had an invalid Type field.
Discarded: Config BPDU length	The number of Configuration BPDUs with incorrect length.
Discarded: TCN BPDU length	The number of Topology Change Notification BPDUs that had an incorrect length.

Enabling one or more STP debugging modes for a period of time displays information for STP troubleshooting (Table 18) to the port on which the switch received the command, or to the console.

**Table 18: STP debugging options**

Option	Debug Mode	Description
MSG	Message	Decoded display of received and transmitted STP packets
PKT	Packet	Raw ASCII display of received and transmitted STP packets
STATE	State	Port state transitions
ALL	All	All debug options

To enable, disable or show the debug modes, use the commands:

```
ENABLE STP DEBUG={MSG|PKT|STATE|ALL} PORT={port-list|ALL}
[OUTPUT=CONSOLE] [TIMEOUT={1..4000000000|NONE}]

DISABLE STP DEBUG={MSG|PKT|STATE|ALL} PORT={port-list|ALL}

SHOW STP DEBUG
```

Set OUTPUT to CONSOLE if using this command in a script. Each of the debug modes can be enabled or disabled independently. Use the TIMEOUT parameter to prevent the switch or display from being overloaded with debugging data.

**Figure 18: Example output from the SHOW STP DEBUG command**

Port	Enabled Debug Modes	Output	Timeout
Port1	MSG, PKT, STATE	16	42
Port2	STATE	16	12345
Port3	None		

**Table 19: Parameters displayed in the output of the SHOW STP DEBUG command**

Parameter	Meaning
Port	The port number on the switch.
Enabled Debug Modes	The debugging option for the port; one of "MSG", "PKT", "STATE" or "NONE".
Output	The output device for the port.
Timeout	The time in seconds during which the port will stay in debug mode.

If necessary, all the STP configuration that users have created on the switch can be removed, so that all STPs except the default STP are destroyed, and all other defaults are restored, using the command:

```
PURGE STP
```



**The PURGE STP command should be used with caution, and generally only before major reconfiguration of the switch, as it removes all STP configuration entered on the switch.**

## IGMP Snooping

IGMP (*Internet Group Management Protocol*) is used by IP hosts to report their multicast group memberships to routers and switches. IP hosts join a multicast group to receive broadcast messages directed to the multicast group address. IGMP is an IP-based protocol and uses IP addresses to identify both the multicast groups and the host members. For a VLAN-aware devices, this means multicast group membership is on a per-VLAN basis. If at least one port in the VLAN is a member of a multicast group, by default multicast packets will be flooded onto all ports in the VLAN.

*IGMP snooping* enables the switch to forward multicast traffic intelligently on the switch. The switch listens to IGMP membership reports, queries and leave messages to identify the switch ports that are members of multicast groups. Multicast traffic will only be forwarded to ports identified as members of the specific multicast group.

IGMP is enabled using the command:

```
ENABLE IP IGMP
```

IGMP snooping is then enabled on a VLAN using the command:

```
ENABLE IP IGMP INTERFACE={VLAN-vlanname|VLANvid}
```

The switch will snoop IGMP packets transiting the VLAN and only forward multicast packets to the ports which have seen a membership report from network devices connected to those ports, instead of being forwarded to all ports belonging to the VLAN.

The command:

```
SET IP IGMP TIMEOUT=1..65535 QUERYINTERVAL=1..65535
```

sets operational parameters for IGMP. The QUERYINTERVAL parameter specifies the time interval, in seconds, at which IGMP Host Membership Queries are sent if this switch is elected the designated router for the LAN. The default is 125.

The TIMEOUT parameter specifies the longest interval, in seconds, that a group will remain in the local group database without receiving a Host Membership Report. The default is 270. If a value is specified for QUERYINTERVAL without specifying a value for TIMEOUT, TIMEOUT is calculated as  $2 \times (\text{QUERYINTERVAL} + 10)$ . The 10 seconds is the variation that hosts use when sending Host Membership Reports. If a timeout interval is specified, it will override any calculated value.

The command:

```
SHOW IP IGMP
```

displays information about IGMP, IGMP snooping, and multicast group membership for each VLAN-based IP interface (Figure 19 on page 49, Table 20 on page 49).

**Figure 19: Example output from the SHOW IP IGMP command.**

```

IGMP Protocol
-----
Status ..... Enabled
Default Query Interval ..... 125 secs
Default Timeout Interval ..... 270 secs

Interface Name ..... vlan10                (DR)
Group List .....
  Group. 224.0.1.17          Last Adv. 192.168.1.130      Refresh time 27
  Group. 224.0.1.43          Last Adv. 192.168.1.130      Refresh time 27
  Group. 224.0.1.66          Last Adv. 192.168.1.140      Refresh time 27
-----

```

**Table 20: Parameters in the output of the SHOW IP IGMP command.**

Parameter	Meaning
Status	The status of IGMP; one of "Enabled" or "Disabled".
Default Query Interval	The interval at which Host Membership Queries are sent.
Default Timeout Interval	The interval after which entries will be removed from the group database, if no Host Membership Report is received.

**Table 20: Parameters in the output of the SHOW IP IGMP command. (Continued)**

Parameter	Meaning
Interface Name	The name of an IP interface, followed by "(DR)" if the interface is acting as the designated router.
Group List	A list of multicast group memberships for this interface.
Group.	The group multicast address.
Last Adv.	The last host to advertise the membership report.
Refresh time	The time interval (in seconds) until the membership group will be deleted if it does not receive another membership report before then.

## Triggers

The Trigger Facility can be used to automatically run specified command scripts when particular triggers are activated. When a trigger is activated by an event, global parameters and parameters specific to the event are passed to the script that is run. For a full description of the Trigger Facility, see the *Trigger Facility* chapter in the *AR Series Router Reference Manual* at <http://www.alliedtelesyn.co.nz/support/rapier/>.

The switch can generate triggers to activate scripts when a fibre uplink port loses or gains coherent light. To create or modify a switch trigger, use the commands:

```
CREATE TRIGGER=trigger-id MODULE=SWITCH
  EVENT={LIGHTOFF|LIGHTON} PORT=port [AFTER=hh:mm]
  [BEFORE=hh:mm] [DATE=date|DAYS=day-list] [NAME=name]
  [REPEAT={YES|NO|ONCE|FOREVER|count}] [SCRIPT=filename...]
  [STATE={ENABLED|DISABLED}] [TEST={YES|NO|ON|OFF}]

SET TRIGGER=trigger-id PORTS={port-list|ALL} [AFTER=hh:mm]
  [BEFORE=hh:mm] [DATE=date|DAYS=day-list] [NAME=name]
  [REPEAT={YES|NO|ONCE|FOREVER|count}]
  [TEST={YES|NO|ON|OFF}]
```

The following sections list the events that may be specified for the EVENT parameter, the parameters that may be specified as *module-specific-parameters*, and the arguments passed to the script activated by the trigger.

<b>Event</b>	LIGHTOFF
<b>Description</b>	The fibre port specified by the PORT parameter has just lost coherent light.
<b>Parameters</b>	The following command parameter(s) must be specified in the CREATE/SET TRIGGER commands:

Parameter	Description
PORT= <i>port</i>	The port on which the event will activate the trigger.

**Script Parameters** The trigger passes the following parameter(s) to the script:

Argument	Description
%1	The port number of the port which has just lost coherent light.

<b>Event</b>	LIGHTON
<b>Description</b>	The fibre port specified by the PORT parameter has just gained coherent light.
<b>Parameters</b>	The following command parameter(s) must be specified in the CREATE/SET TRIGGER commands:

Parameter	Description
PORT= <i>port</i>	The port on which the event will activate the trigger.

**Script Parameters** The trigger passes the following parameter(s) to the script:

Argument	Description
%1	The port number of the port which has just gained coherent light.

## Layer 3 Switching

The Rapier switch provides Layer 3 switching and routing over VLANs. Once a VLAN has been created (see “Virtual LANs” on page 25), the VLAN name can be used wherever a logical interface is required in commands for configuring routing protocols.

VLAN names are of the form:

VLAN-*vlannname*

or

VLAN*n*

where *vlannname* is the manager-assigned name of the VLAN, and *n* is the VLAN identifier (VID).

For example, if a VLAN is created using the command:

```
CREATE VLAN=admin VID=11
```

then the following names can be used to identify the VLAN in routing commands:

```
vlan-admin
```

```
vlan11
```

The following sections illustrate the use of VLANs for IP, RIP, IPX, AppleTalk and RSVP. For a complete description of all the protocols supported by the switch, see the AR Series Router Reference Manual at <http://www.alliedtelesyn.co.nz/support/rapier/>.

## Internet Protocol (IP)

The switch performs IP routing at wire speed between VLANs. To add the admin VLAN as an IP interface, use either of the following commands:

```
ADD IP INTERFACE=vlan-admin
```

```
ADD IP INTERFACE=vlan11
```

The command:

```
SHOW IP INTERFACE
```

displays the interfaces enabled for IP routing (Figure 20 on page 52).

Figure 20: Example output from the SHOW IP INTERFACE command.

Interface Pri. Filt	Type Pol.Filt	IP Address Network Mask	Bc Fr MTU VJC	PArp GRE	Filt OSPF Met.	RIP Met. DBcast	SAMode Mul.	IPSc
LOCAL	-	Not Set	- n	-	---	-	-	--
---	----	-	-	-	---	-	-	---
vlan11	Static	192.168.163.39	1 y	On	---	01	Pass	--
---	---	255.255.255.0	1500	-	---	0000000001	No	On
ppp1	Dynamic	0.0.0.0	1 y	-	---	01	Pass	--
---	---	255.255.255.255	1500	Off	---	0000000001	No	On
-----	-----	-----	-----	-----	-----	-----	-----	-----

## Routing Information Protocol (RIP)

Routing protocols such as RIPv1 and RIPv2 can be enabled on a VLAN. For example, the command:

```
ADD IP RIP INTERFACE=vlan11 SEND=RIP2 RECEIVE=BOTH
```

enables RIPv2 on the admin VLAN. The command:

```
SHOW IP RIP
```

displays information about RIP (Figure 21 on page 52).

Figure 21: Example output from the SHOW IP RIP command.

Interface	Circuit/DLCI	IP Address	Send	Receive	Demand	Auth	Password
-----	-----	-----	-----	-----	-----	-----	-----
vlan11	-	-	RIP2	BOTH	NO	NO	
ppp0	-	172.16.249.34	RIP1	RIP2	YES	PASS	*****
-----	-----	-----	-----	-----	-----	-----	-----

## Novell IPX

The switch's implementation of the Novell IPX protocol uses the term *circuit* to refer to a logical connection over an *interface*, similar to an X.25 permanent virtual circuit (PVC) or a Frame Relay Data Link Connection (DLC). The term *interface* is used to refer to the underlying physical interface, such as VLAN, Ethernet, Point-to-Point (PPP) and Frame Relay.

To create IPX circuit 1 with the Novell network number 129 over the admin VLAN, use the command:

```
ADD IPX CIRC=1 INTERFACE=vlan11 NETWORK=129 ENCAP=802.3
```

The command:

```
SHOW IPX CIRCUIT
```

displays information about the circuits configured for IPX (Figure 22 on page 53).

**Figure 22: Example output from the SHOW IPX CIRCUIT command.**

```

IPX CIRCUIT information

Name ..... Circuit 1
Status ..... enabled
Interface ..... vlan11   (802.3)
Network number ..... c0e7230f
Station number ..... 0000cd000d26
Link state ..... up
Cost in Novell ticks ..... 1
Type20 packets allowed ..... no
On demand ..... no

Spoofing information
Keep alive spoofing ..... no
SPX watch dog spoofing ..... no
On SPX connection failure .... UPLINK
On end of SPX spoofing ..... UPLINK

RIP broadcast information
Change broadcasts ..... yes
General broadcasts ..... yes
General broadcast interval ... 60 seconds
Maximum age ..... 180 seconds

SAP broadcast information
Change broadcasts ..... yes
General broadcasts ..... yes
General broadcast interval ... 60 seconds
Maximum age ..... 180 seconds

Filter information
Filters ..... none

```

## AppleTalk

To create an AppleTalk port (interface) associated with the admin VLAN, use the command:

```
ADD APPLE PORT INTERFACE=vlan11
```

The command:

```
SHOW APPLE PORT
```

displays information about the ports configured for AppleTalk (Figure 23 on page 54).

**Figure 23: Example output from the SHOW APPLE PORT command.**

```

Appletalk Port Details
-----
Port Number ..... 1
Interface ..... vlan11
ifIndex ..... 1
Node ID ..... 217
Network Number ..... 22
Network Range Start ..... 22
Network Range End ..... 22
State ..... ACTIVE
Seed ..... NO
Seed Network Start ..... 0
Seed Network End ..... 0
Hint ..... YES
Hint Node ID ..... 179
Hint Network ..... 22
Default Zone ..... -

Zone List is Empty
-----

```

## Resource Reservation Protocol (RSVP)

The Resource Reservation Protocol (RSVP) enables the receiver of a traffic flow to make the resource reservations necessary to ensure that the receiver obtains the desired QoS for the traffic flow.

RSVP is disabled by default. To enable RSVP, use the command:

```
ENABLE RSVP
```

Each IP interface that is to receive and process RSVP messages and accept reservation requests must be enabled. To enable RSVP on the admin VLAN, use the command:

```
ENABLE RSVP INTERFACE=vlan11
```

The command:

```
SHOW RSVP INTERFACE
```

displays information about the interfaces enabled for RSVP (Figure 24 on page 54).

**Figure 24: Example output from the SHOW RSVP INTERFACE command.**

```

RSVP Interfaces

```

Interface	Enabled	Maximum Bandwidth(%)	Reserved Bandwidth(%)	No. Of Reservations	Debug	Encap
Dynamic	No	75	0	0	None	RAW
vlan11	Yes	75	0	1	None	RAW
ppp0	Yes	75	0	0	None	RAW

## Layer 3 LAN/WAN Routing

In addition to Layer 2 and Layer 3 switching, the Rapier switch implements almost all of the AR routing software suite from Allied Telesyn's AR series of routers, providing a wide array of multiprotocol routing, security and network management features.



*IP routing is performed at wire-speed. Other Layer 3 routing is performed by the CPU, and increasing the routing load on the CPU decrease its performance.*

Some features require the addition of WAN interfaces via NSMs and PICs installed in the NSM bay on the rear of the switch.

Features provided by the AR routing software suite include:

- IP routing
- Network Address Translation (NAT) (not between switch ports)
- Dynamic IP Address Assignment
- IP Dynamic Filtering Firewall
- IP Multihoming
- OSPF (Open Shortest Path First) (not between switch ports)
- RIP and RIPv2
- DNS Relay
- Demand IP
- IP Filtering (not between switch ports)
- IP Packet Prioritisation (not between switch ports)
- Generic Routing Encapsulation (GRE)
- Basic Rate and Primary Rate access to Integrated Services Digital Network (ISDN) services, with dial-on-demand and channel aggregation.
- Time Division Multiplexing (TDM) over G.703 links
- Frame Relay
- X.25
- ARP, Proxy ARP and Inverse ARP address resolution protocols.
- IPX routing
- Demand IPX
- Demand IP and IPX
- IPX/SPX Spoofing
- IPX Filtering (not between switch ports)
- AppleTalk routing
- BACP (Bandwidth Allocation Control Protocol)
- PPP Multilink
- PPP over Ethernet (PPPoE)
- Bandwidth on Demand
- CLI, PAP and CHAP

- Virtual Router Redundancy Protocol (VRRP) for fault tolerant internet gateways (on NSM ports only)
- IPsec
- ISAKMP Key Management
- Data Compression
- Predictor Data Compression
- STAC Data Compression
- Nemesis Stateful Inspection Firewall
- SecureShell Remote Management
- Resource Reservation Protocol (RSVP)
- L2TP
- Telnet client and server.
- A sophisticated and configurable event logging facility for monitoring and alarm notification to single or multiple management centres.
- Triggers for automatic and timed execution of commands in response to events.
- Scripting for automated configuration and centralised management of configurations.
- Dynamic Host Configuration Protocol (DHCP) for automatically assigning IP addresses and other configuration information to PCs and other hosts on TCP/IP networks.
- Support for the Simple Network Management Protocol (SNMP), standard MIBs and the Allied Telesyn Enterprise MIB, enabling the switch to be managed by a separate SNMP management station.
- An HTTP client that allows files to be downloaded directly from a web server to the switch's FLASH memory, and an HTTP server that serves web pages from FLASH.

For a complete description of the AR router software suite, see the AR Series Router Reference Manual (Document Number C613-03016-00 Rev B) at <http://www.alliedtelesyn.co.nz/support/rapier/>.

## SNMP and MIBs

The switch's implementation of SNMP is based on RFC 1157 "A Simple Network Management Protocol (SNMP)", and RFC 1812, "Requirements for IP Version 4 Routers". The SNMP agent is disabled by default. To enable SNMP, use the command:

```
ENABLE SNMP
```

SNMP *communities* are the main configuration item in the switch's SNMP agent, and are defined in terms of a list of IP addresses which define the SNMP application entities (trap hosts and management stations) in the community. An SNMP community is created using the command:

```
CREATE SNMP COMMUNITY=name [ACCESS={ READ | WRITE } ]
[TRAPHOST=ipadd] [MANAGER=ipadd]
[OPEN={ ON | OFF | YES | NO | TRUE | FALSE } ]
```

Authentication failure traps and link state traps can be enabled using the commands:

```
ENABLE SNMP AUTHENTICATE_TRAP
ENABLE INTERFACE=interface LINKTRAP
```

where *interface* is the name of an interface, such as vlan11.

The command:

```
SHOW SNMP
```

displays the current state and configuration of the SNMP agent (Figure 25 on page 57).

**Figure 25: Example output from the SHOW RSVP INTERFACE command.**

```
SNMP configuration:
Status ..... Enabled
Authentication failure traps .... Enabled
Community ..... public
Access ..... read-only
Status ..... Enabled
Traps ..... Enabled
Open access ..... Yes
Community ..... Administration
Access ..... read-write
Status ..... Disabled
Traps ..... Disabled
Open access ..... No

SNMP counters:
inPkts ..... 0          outPkts ..... 0
inBadVersions ..... 0    outTooBigs ..... 0
inBadCommunityNames ..... 0 outNoSuchNames ..... 0
inBadCommunityUses ..... 0 outBadValues ..... 0
inASNParseErrs ..... 0    outGenErrs ..... 0
inTooBigs ..... 0         outGetRequests ..... 0
inNoSuchNames ..... 0     outGetNexts ..... 0
inBadValues ..... 0       outSetRequests ..... 0
inReadOnly ..... 0        outGetResponses ..... 0
inGenErrs ..... 0         outTraps ..... 0
inTotalReqVars ..... 0
inTotalSetVars ..... 0
inGetRequests ..... 0
inGetNexts ..... 0
inSetRequests ..... 0
inGetResponses ..... 0
inTraps ..... 0
```

The following MIBs are supported in the field trial release:

- MIB II (RFC 1213)
- Ethernet MIB (RFC 1643)
- Trap MIB (RFC 1215)
- RMON Groups 1, 2, 3, and 9 (RFC 1757)
- AR Router portion of the ATI/ATKK Enterprise MIB
- Portions of the Extended Interface MIB (RFC 1573)

Additional MIBs will be added in future releases.

## Availability

---

The Rapier 24 switch is supplied with Software Release 2.1.0 installed.

During the field trial, software upgrades and upgrade information for the Rapier 24 will be available from the Allied Telesyn Research web site <http://www.alliedtelesyn.co.nz/support/rapier/>.

The LOAD command can be used to download software upgrades directly from the Allied Telesyn Research web site to the switch's FLASH memory. Use the SET INSTALL command to enable the new software release (*Example: Install Software Upgrade for Rapier Switch* on page 16).